

XTREME-R: Towards More Challenging and Nuanced Multilingual Evaluation

Sebastian Ruder¹, Noah Constant², Jan Botha², Aditya Siddhant², Orhan Firat², Jinlan Fu³, Pengfei Liu⁴, Junjie Hu⁴, Dan Garrette², Graham Neubig⁴, Melvin Johnson²

¹DeepMind ²Google Research ³Fudan University ⁴Carnegie Mellon University

Abstract

Machine learning has brought striking advances in multilingual natural language processing capabilities over the past year. For example, the latest techniques have improved the state-of-the-art performance on the XTREME multilingual benchmark by more than 13 points. While a sizeable gap to human-level performance remains, improvements have been easier to achieve in some tasks than in others. This paper analyzes the current state of cross-lingual transfer learning and summarizes some lessons learned. In order to catalyze meaningful progress, we extend XTREME to XTREME-R, which consists of an improved set of ten natural language understanding tasks, including challenging language-agnostic retrieval tasks, and covers 50 typologically diverse languages. In addition, we provide a massively multilingual diagnostic suite (MULTICHECKLIST) and fine-grained multi-dataset evaluation capabilities through an interactive public leaderboard to gain a better understanding of such models.

1 Introduction

Most research in natural language processing (NLP) to date has focused on developing methods that work well for English and a small set of other high-resource languages (Joshi et al., 2020). In contrast, methods for other languages can be vastly more beneficial as they enable access to language technology for more than three billion speakers of low-resource languages and prevent the NLP community from overfitting to English. Motivated by these benefits, the area of multilingual NLP has attracted increasing interest recently.

However, evaluating multilingual models is challenging as it requires assessing performance on a wide range of typologically distinct languages in the face of limited heterogeneous data sources. Recently large-scale benchmarks such as XTREME (Hu et al., 2020) and XGLUE (Liang et al., 2020)

	XTREME	XTREME-R
# of languages	40	50
# of tasks	9	− 2 + 3 = 10
Task categories	Classification, structured prediction, QA, retrieval	+language-agnostic retrieval
Analysis tools	—	MULTICHECKLIST, Explainaboard
Leaderboard	Static	Interactive, +metadata

Table 1: Overview of XTREME and XTREME-R.

have been introduced, consolidating existing multilingual tasks and covering tens of languages. When XTREME was released, the gap between the best-performing baseline, XLM-R Large (Conneau et al., 2020), and human-level performance was roughly 25. This has since shrunk to less than 12 points, a much smaller but still substantial gap compared to the difference from human-level performance observed in English transfer learning (Wang et al., 2019a), which has recently been closed entirely on some evaluation suites (He et al., 2021).

In order to examine the nature of this progress, we first perform an analysis of state-of-the-art multilingual models on XTREME. We observe that progress has not been uniform, but concentrated on cross-lingual retrieval tasks where fine-tuning on other tasks and pre-training with parallel data lead to large gains. On other task categories improvements are more modest. Models still generally perform poorly on languages with limited data and non-Latin scripts. Fine-tuning on additional translated data generally leads to the best performance.

Based on this analysis, we propose XTREME-R (XTREME Revisited), a new benchmark with the dual purpose of ensuring that research in multilingual NLP focuses on the most challenging problems and equipping researchers with a broader set of tools to better understand their models (see Table 1 for a brief overview). XTREME-R follows in its predecessor’s footsteps by being massively multilingual, diverse, and accessible. It expands on

XTREME by covering 50 typologically diverse languages and 10 challenging, diverse tasks. To make retrieval more difficult, we introduce two new tasks that focus on “language-agnostic” retrieval (Roy et al., 2020), where targets must be retrieved from a large multilingual candidate pool. We additionally establish new state-of-the-art mT5 (Xue et al., 2021) and translate-train baselines for our tasks.

XTREME-R aims to move away from a single aggregate metric summarizing a model’s performance and towards a more nuanced evaluation and comparison of multilingual models (Ethayarajh and Jurafsky, 2020; Linzen, 2020). To this end, we introduce an extensible multilingual diagnostic and evaluation suite that consists of two main components: a) MULTICHECKLIST, a test suite (Ribeiro et al., 2020) for probing question answering capabilities in 50 languages. This test suite is the first of its kind and enables direct evaluation of fine-grained capabilities in a massively multilingual setting. b) We extend the multi-dataset evaluation framework EXPLAINBOARD (Fu et al., 2020; Liu et al., 2021) to additional tasks and the multilingual setting. This framework allows us to break down performance based on language and task-specific attributes, which enables a more nuanced diagnosis of a model’s behaviour.

We also make several logistic improvements to improve XTREME-R’s utility as a leaderboard. To make it easier to choose the best model for a use case, each submission is required to provide metadata such as the number of parameters and amount of pre-training data, which we make available via an interactive leaderboard. We also introduce task and language-specific sub-leaderboards to invite submissions of dedicated models.

In sum, we make the following contributions: a) an analysis of progress in cross-lingual modeling; b) an improved benchmark covering 50 languages, including a newly created retrieval task (Mewslin); c) a massively multilingual diagnostic suite; d) fine-grained evaluation capabilities; e) experiments and analyses of state-of-the-art models; and f) an interactive metadata-rich leaderboard.

2 Examining the State of Multilingual Benchmarking

2.1 Background

Benchmarking is critical to evaluate general-purpose language understanding technologies. To this end, benchmarks like GLUE (Wang et al.,

2019b) and SuperGLUE (Wang et al., 2019a) provide a way to assess the transfer learning capabilities of various models. However, these benchmarks focus only on English. On the other hand, cross-lingual approaches have been evaluated on a wide range of disparate tasks (Hu et al., 2020). XTREME was proposed as a platform to unify this fragmented evaluation landscape and to catalyze advances in cross-lingual learning by including a diverse set of tasks and languages. It consists of 9 tasks covering 40 diverse languages, which can be grouped into 4 broad task types (see §3.1 for details): classification (XNLI, PAWS-X), structured prediction (UD-POS, WikiANN-NER), question answering (XQuAD, MLQA, TyDiQA-GoldP), and retrieval (Tatoeba, BUCC). XTREME focuses on *zero-shot cross-lingual transfer*, i.e. models can be pre-trained on any multilingual data and are fine-tuned only in English. Similarly, XGLUE (Liang et al., 2020), another cross-lingual benchmark focuses on a smaller number of less typologically diverse languages, and includes generation tasks. Other non-English benchmarks focus on specific linguistic phenomena, e.g. code-switching (Khanuja et al., 2020); languages, e.g. Indonesian (Willie et al., 2020) and Persian (Khashabi et al., 2020); and language families, e.g. Indian languages (Kakwani et al., 2020).

2.2 An Analysis of XTREME

As of April 15, 2021, all submissions to the XTREME leaderboard are large-scale Transformers (Vaswani et al., 2017) trained with masked language modeling (MLM; see Appendix A for further details). We analyze the performance of these models on the XTREME leaderboard in Figure 1.¹ Overall, multilingual models have improved the average performance on XTREME from 55.8 to 81.4. Much of this improvement is concentrated in the retrieval-based tasks where performance increased from 47.7 (mBERT) to 92.7 (VECO). In contrast, performance on question answering and structured prediction tasks has improved only slightly.

Breaking down performance by language family, on Tatoeba (Figure 1c) recent models still struggle with a few low-resource languages. Models perform well for most other languages and their scores are concentrated in a relatively small range. On MLQA (Figure 1b), scores have increased slightly

¹This evaluation compares “models + data” as models employ different types of data during training. We document this information in Appendix H.

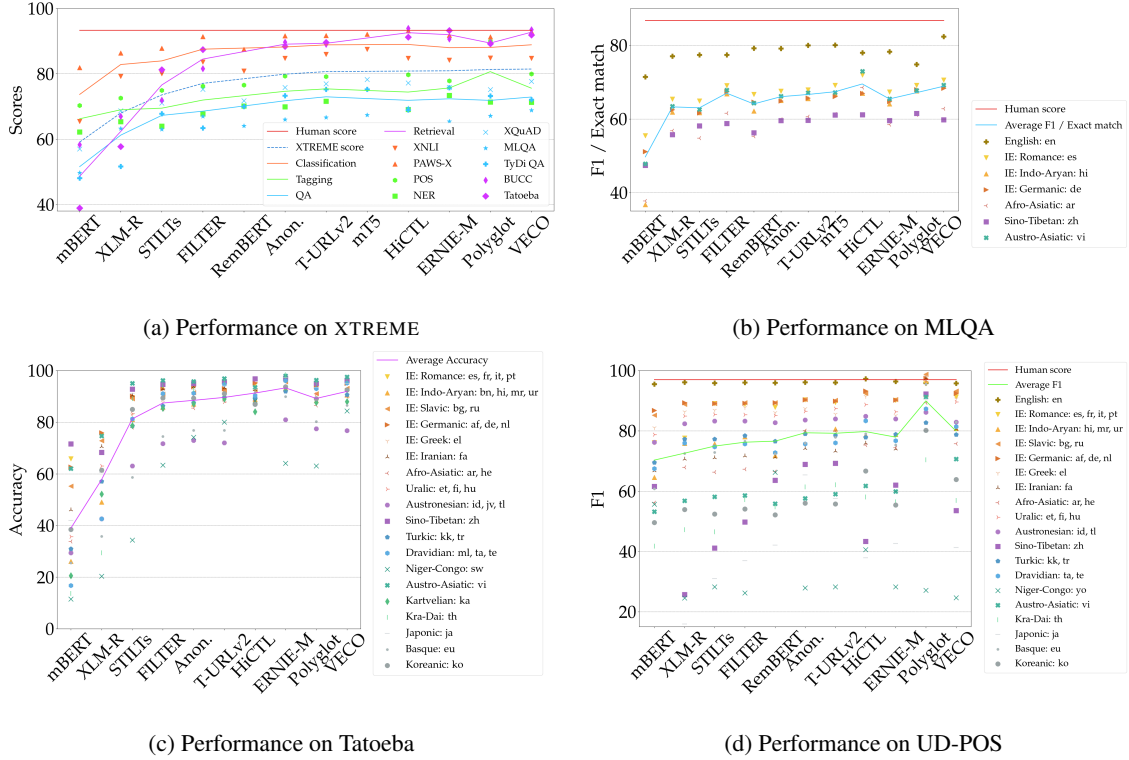


Figure 1: Performance of models (a) on the XTREME leaderboard across all nine XTREME tasks, (b) on the MLQA question answering dataset, (c) on the Tatoeba retrieval task, and (d) on Universal Dependencies POS tagging across language families. Models are ordered based on their XTREME score (a). Results of models that do not evaluate on a task category are omitted, i.e. RemBERT for retrieval and mT5 for retrieval and tagging.

but remain well below performance on English. On POS tagging (Figure 1d), scores remain largely the same; performance is lower for some languages with non-Latin scripts and low-resource languages. We show the scores for the remaining tasks in Appendix B. The remaining gap to English performance on these tasks is partially an artefact of the evaluation setup: zero-shot cross-lingual transfer from English favors English representations whereas models fine-tuned on in-language monolingual data perform more similarly across languages (Clark et al., 2020; Hu et al., 2020).

Overall, representations from *token-level* MLM pre-training are of limited use for cross-lingual *sentence* retrieval, as evidenced by the comparatively poor performance of the mBERT and XLM-R models. Fine-tuning on sentence-level tasks (Phang et al., 2020; Fang et al., 2021) can mitigate this. The strong performance of recent models such as VECO and ERNIE-M on the retrieval tasks can be attributed to a combination of parallel data and new pre-training objectives that make use of it. Pre-training on parallel data improves performance on retrieval by making the pre-training task more sim-

ilar to the downstream setting but does not significantly improve performance on other tasks. *Fine-tuning* on automatically translated task-specific data yields strong gains and is used by most recent models to achieve the best performance (Hu et al., 2020; Ouyang et al., 2020; Luo et al., 2020). Nevertheless, key challenges such as how to learn robust cross-lingual syntactic and semantic processing capabilities during *pre-training* remain.

3 XTREME-R

In order to encourage the NLP community to tackle challenging research directions in pursuit of better cross-lingual model generalization, we propose XTREME-R (XTREME Revisited). XTREME-R shares its predecessor’s core design principles for creating an accessible benchmark to evaluate cross-lingual transfer but makes some key changes.

First, XTREME-R focuses on the tasks that have proven to be hardest for current multilingual models. To this end, it drops XTREME’s PAWS-X and BUCC tasks since recent advances have left less room for further improvement, and they cover only a small number of less diverse languages. They

Task category	Task	Train	Dev	Test	Test sets	Lang.	Task	Metric	Domain
Classification	XNLI	392,702	2,490	5,010	translations	15	NLI	Accuracy	Misc.
	XCOPA	33,410+400	100	500	translations	11	Reasoning	Accuracy	Misc.
Struct. prediction	UD-POS	21,253	3,974	47-20,436	ind. annot.	37 (104)	POS	F1	Misc.
	WikiANN-NER	20,000	10,000	1,000-10,000	ind. annot.	47 (176)	NER	F1	Wikipedia
QA	XQuAD	87,599	10,570	1,190	translations	11	Span extraction	F1 / EM	Wikipedia
	MLQA			4,517-11,590	translations	7	Span extraction	F1 / EM	Wikipedia
	TyDiQA-GoldP	3,696	634	323-2,719	ind. annot.	9	Span extraction	F1 / EM	Wikipedia
Retrieval	Tatoeba	87,599	10,570	1,000	translations	38 (122)	Sentence retrieval	Accuracy	Misc.
	Mewsl-X	116,903	10,252	428-1,482	ind. annot.	11 (50)	Lang. agn. retrieval	mAP@20	News
	LAReQA XQuAD-R	87,599	10,570	1,190	translations	11	Lang. agn. retrieval	mAP@20	Wikipedia

Table 2: The tasks in XTREME-R. For tasks that have training and dev sets in other languages, we only report the English numbers. We report the number of test examples per target language, the nature of the test sets (whether they are translations of English data or independently annotated), and the size of the intersection with our selected languages (the total number of languages is in brackets). For WikiANN-NER and UD-POS, sizes are in sentences. XCOPA includes training data of SIQA (Sap et al., 2019) and Tatoeba uses SQuAD v1.1 data for fine-tuning.

are replaced instead by three new, more challenging tasks: one focusing on causal commonsense reasoning (§3.2.1) and two focusing on harder retrieval scenarios (§3.2.2), as this has been the task category where gains have been easiest to achieve. We retain XTREME’s seven other tasks as each still presents substantial challenges for state-of-the-art cross-lingual models (§3.1). Overall, XTREME-R includes 10 diverse tasks, summarized in Table 2.

We also make changes to the structured prediction tasks, NER and POS. Instead of only providing examples as lists of tokens, XTREME-R always provides the full text of an input sentence, thus ensuring that the entire benchmark now supports research on models that operate directly from the raw input string (Clark et al., 2021). Furthermore, XTREME-R adopts a more realistic version of the NER task in which no gold tokenization is provided at all, meaning that systems will either have to use model-predicted tokens or embrace tokenization-free approaches. Finally, XTREME-R provides a multilingual diagnostic and evaluation suite (§3.4).

3.1 Retained Tasks

We retain the XNLI (Conneau et al., 2018), UD-POS (Nivre et al., 2018), WikiANN-NER (Pan et al., 2017), XQuAD (Artetxe et al., 2020a), MLQA (Lewis et al., 2020), TyDiQA-GoldP (Clark et al., 2020), and Tatoeba (Artetxe and Schwenk, 2019) tasks from XTREME (see Appendix C).

3.2 New Tasks

3.2.1 Multilingual Causal Reasoning

XCOPA The Cross-lingual Choice of Plausible Alternatives (Ponti et al., 2020) dataset asks models to decide which of two sentences causally follows

a premise sentence. The XCOPA authors translated and re-annotated the validation and test sets of the English COPA (Roemmele et al., 2011) dataset into 11 languages, which we use for evaluation. The English COPA training set together with the Social IQa (Sap et al., 2019) training data are used for training. While accuracy on the English COPA recently reached 94.8% (Raffel et al., 2020), the state-of-the-art on XCOPA is only around 70%.

3.2.2 Retrieval from a Multilingual Pool

Many previous retrieval benchmarks assume that the entire candidate pool is in a single language. For instance, a French query will be used to search over only English candidates. However, practical settings often violate this assumption, e.g. the answer to a question may be available in any number of languages, possibly different from the query language. Models that cannot compare the appropriateness of retrieval results across languages are thus ineffective in such real-world scenarios.

XTREME-R includes two new related cross-lingual retrieval tasks. The first seeks to measure the extent to which cross-lingual representations are “strongly aligned” (Roy et al., 2020), i.e. they place the semantically most related text pairs (e.g. a question and its answer) closest together in representation space, *regardless* of their language identities. The second analogously frames entity linking as retrieving from a multilingual pool of entity descriptions, given an entity mention in context (Botha et al., 2020). For both, we report performance as mean average precision at 20 (mAP@20). **LAReQA** Language Agnostic Retrieval Question Answering (Roy et al., 2020) is a sentence retrieval task. Each query has target answers in multiple languages, and models are expected to rank all correct

Test	Template	Generated test
Comparisons	{first_name} is {adj[0]} than {first_name1}. Who is less {adj[1]}?	C: Ben is smaller than Frank. Q: Who is less small?
Intensifiers	{first_name} {state} {very} בקשר לפרויקט. מי הכי פחות {state} בקשר לפרויקט?	C: עמנאל שמח בקשר לפרויקט. יצחק שמח ביותר בקשר לפרויקט. Q: מי הכי פחות שמח בקשר לפרויקט?
Properties	{attribute1} و {attribute2}. يوجد ورق حائط في الغرفة. {obj[0]} في الغرفة. {obj[1]} أي {obj[1]} هو {property2}؟	C: يوجد ورق حائط في الغرفة. ورق الحائط هو ضئيل ومربع. Q: أي شكل هو ورق الحائط؟
Job vs Nationality	{first_name} একজন {profession} এবং {nationality}। {first_name} এর জাতীয়তা কী?	C: হালিম একজন ওয়েবস্টেস এবং চীনা। Q: হালিম এর জাতীয়তা কী?

Table 3: CHECKLIST templates and generated tests for different capabilities in English, Hebrew, Arabic, and Bengali. Words in curly brackets { . . . } are placeholders; see Ribeiro et al. (2020) for more information.

answers above all incorrect answers, regardless of language. We use the LARQA XQuAD-R dataset which contains 13,090 questions each of which has 11 target answers (in 11 distinct languages) within the pool of 13,014 candidate answer sentences. Following Roy et al. (2020), we fine-tune models on the SQuAD v1.1 train set. The fine-tuned model is used to rank the 13K candidates for each question. **Mewsl-X** Mewsl (Multilingual Entities in News, linked) is an automatically extracted dataset that requires linking a contextual entity mention to its entry in a language-agnostic knowledge base by retrieving the entity’s description from a multilingual candidate pool (Botha et al., 2020). For XTREME-R, we derive Mewsl-X as a new variant of Mewsl-9, still linking against WikiData (Vrandečić and Krötzsch, 2014). Mewsl-X features 15K mentions in 11 languages: given a mention in context, the task is to retrieve the single correct target entity description from a candidate pool ranging over 1M candidates across all 50 languages of XTREME-R. Fine-tuning is done on a predefined set of English-only mention-entity pairs randomly sampled from English Wikipedia hyperlinks (see Appendix E for further details).

For our baseline systems on both tasks, we follow previous work (Roy et al., 2020; Botha et al., 2020) and train a dual encoder initialized from the pre-trained model weights, optimizing for an in-batch sampled softmax loss (Gillick et al., 2018).

3.3 Languages

XTREME-R adds the following ten languages to XTREME: Haitian Creole, Cusco Quechuan, Wolof, Lithuanian, Punjabi, Gujarati, Polish, Ukrainian, Azerbaijani, and Romanian. In total, XTREME-R covers the following 50 languages (shown using their ISO 639-1 codes for brevity; new languages are bolded) belonging to 14 language families and two isolates: af, ar, **az**, bg, bn, de, el, en, es, et, eu, fa, fi, fr, **gu**, he, hi, **ht**, hu, id, it, ja, jv, ka, kk, ko,

lt, ml, mr, ms, my, nl, **pa**, **pl**, pt, **qu**, **ro**, ru, sw, ta, te, th, tl, tr, **uk**, ur, vi, **wo**, yo, zh.² XTREME-R is similarly typologically and genealogically diverse as XTREME while covering a larger number of languages (see Appendix D).

3.4 Diagnostic and evaluation suite

To increase the language coverage of low-resource languages in XTREME-R and to enable us to systematically evaluate a model’s cross-lingual generalization ability, we augment XTREME-R with a massively multilingual diagnostic and evaluation suite. Challenge sets and diagnostic suites in NLP (Wang et al., 2019a,b; Belinkov and Glass, 2019) are mostly limited to English, with a few exceptions (Gulordava et al., 2018). As challenge sets are generally created with a human in the loop, the main challenge for creating a large multilingual diagnostic suite is to scale the annotation or translation effort to many languages and to deal with each language’s idiosyncrasies.

MULTICHECKLIST To address this, we build on the CHECKLIST (Ribeiro et al., 2020) framework, which facilitates creating parameterized tests for models. CHECKLIST enables the creation of test cases using templates, which test for specific behavioral capabilities of a model with regard to a downstream task. Importantly, by relying on template-based tests, we can efficiently generate a large number of diverse multilingual test cases by creating a relatively small number of templates in 50 languages.³ We focus on translating English tests, which consist of templates and their fill-in values.⁴ To study the feasibility of creating multilingual test cases at scale, we translate the minimum

²The new languages are covered in both the new tasks as well as in UD-POS, WikiANN-NER, and Tatoeba.

³In contrast, translating an existing test suite or dataset or annotating individual examples for 50 languages would be prohibitively expensive

⁴Templates could alternatively be created by native speakers in each language.

functionality tests (MFT) of CHECKLIST, which probe for general vocabulary and taxonomic knowledge in question answering. We instruct translators to create separate variants of a template to disambiguate linguistic phenomena, such as gender of fill-in values, question type, semantics of properties, etc. We automatically fill names in each language based on data from Wikidata and programmatically consolidate different templates in each language. We show examples of templates and the tests that they generate in different languages in Table 3.

We highlight statistics of the dataset and translation process, instructions to translators, and general challenges of template translation in Appendix F. We believe that parameterized tests are a powerful tool to obtain diverse diagnostics data for otherwise resource-starved languages. We view participatory research (V et al., 2020) with native speakers to create template-based test cases testing for language-specific behaviour as particularly promising.

Multilingual EXPLAINBOARD The standard practice in leaderboards is to average performance across different settings (Wang et al., 2019b,a). While this provides discriminative power, it has limited utility for examining the relative advantages of systems, the characteristics of different datasets and languages, and how these factors relate to each other. To provide more granular evaluation capabilities, we extend Fu et al. (2020); Liu et al. (2021)’s EXPLAINBOARD to the task categories and languages in XTREME-R. EXPLAINBOARD provides a more nuanced impression of a model’s performance on a task by defining task-specific attributes (e.g. *entity length* for NER). The test set is partitioned into different buckets based on the defined attributes and performance is broken down over different attribute values. We define new task-specific attributes for the four task types as well as task-independent attributes (see Appendix K).

Metadata We additionally would like to enable practitioners to rank submissions based on other information. To this end, we ask each submission to XTREME-R for relevant metadata such as the number of parameters, the amount of pre-training data, etc. We will show this information in an interactive leaderboard (see Appendix H for the metadata of current XTREME submissions).

4 Experiments

Training and evaluation setup XTREME-R focuses on *zero-shot cross-lingual transfer* from En-

glish. While recent work (Hu et al., 2020; Lauscher et al., 2020; Hedderich et al., 2020) demonstrates the benefits of fine-tuning on in-language data, we believe the zero-shot scenario remains the most effective way to evaluate the amount of *a priori* multilingual knowledge a pre-trained model captures. Due to variation in cross-lingual evaluation (Keung et al., 2020), we recommend researchers to use the validation set of a *single* target language for development (Artetxe et al., 2020b).

4.1 Baselines

We employ established pre-trained multilingual and models using translations as baselines.

mBERT Multilingual BERT (Devlin et al., 2019) has been pretrained on the Wikipedias of 104 languages using MLM.

XLM-R XLM-R Large (Conneau et al., 2020) uses the same MLM objective with a larger model, and was trained on a magnitude more web data from 100 languages.

mT5 Multilingual T5 (Xue et al., 2021) is an encoder-decoder transformer that frames NLP tasks in a “text-to-text” format. It was pre-trained with MLM on a large multilingual web corpus covering 101 languages. We employ the largest mT5-XXL variant with 13B parameters.

Translate-train To evaluate the impact of MT, we fine-tune mBERT on translations of English training data from Hu et al. (2020). We create new translations for the XCOPA and SIQA data using an in-house MT system.⁵

Translate-train multilingual In addition, we fine-tune both mBERT and mT5 on the combined translated training data of all languages (including the original English data) jointly.

Human performance We use the human performance estimates from XTREME for the retained tasks. For XCOPA we average the proportion of annotated labels disagreeing with the majority label across all languages (Ponti et al., 2020). We are not able to obtain human performance estimates for the new retrieval tasks as identifying a translation among a large number of candidates is too time-consuming for a human to perform.

4.2 Results

We show the main results in Table 4. As in prior work, XLM-R Large generally outperforms mBERT. Fine-tuning helps significantly on Tatoeba

⁵We are unable to produce translations for Quechua.

Model	Avg	Classification		Structured prediction		Question answering			Lang.-agnostic retrieval		Retrieval
		XNLI	XCOPA	UD-POS	WikiANN-NER	XQuAD	MLQA	TyDiQA-GoldP	Mewsli-X	LAReQA	Tatoeba
Metrics		Acc.	Acc.	F1	F1	F1 / EM	F1 / EM	F1 / EM	mAP@20	mAP@20	Acc.
<i>Cross-lingual zero-shot transfer (models are trained on English data)</i>											
mBERT	54.1	66.5	56.1	70.9	62.7	65.1 / 50.4	61.3 / 44.1	58.4 / 43.7	38.6	21.6	43.3
XLm-R Large	65.3	79.2	69.2	75.0	64.4	77.2 / 61.6	72.7 / 54.2	64.3 / 45.8	45.7	40.7	77.3
mT5-XXL	64.6	84.8	74.9	71.5	68.8	81.5 / 66.6	75.6 / 57.3	80.1 / 65.7	41.7*	24.7*	45.7
<i>Translate-train (models are trained on English training data translated to the target language)</i>											
mBERT	-	74.6	57.3	-	-	70.0 / 56.0	65.6 / 48.0	52.4 / 39.5	-	-	-
mBERT, multilingual	-	75.1	57.9	-	-	72.4 / 58.3	67.6 / 49.8	59.5 / 45.8	-	-	-
mT5-XXL, multilingual	-	87.5	77.8	-	-	84.7 / 71.8	76.4 / 58.4	81.9 / 68.6	-	-	-
Human	-	92.8	97.6	97.0	-	91.2 / 82.3	91.2 / 82.3	90.1 / -	-	-	-

Table 4: Overall results of baselines across all XTREME-R tasks. *: Due to compute limitations, mT5-XXL language-agnostic retrieval results are obtained with a frozen rather than a fine-tuned model.

Lang.	Comparisons	Intensifiers	Properties	Job vs Nationality	Animal vs Vehicles	Animal vs Vehicles 2	Avg
fi	10.7	79.8	34.1	4.0	9.2	1.5	23.2
fr	4.0	98.0	15.4	0.8	10.0	13.5	23.6
en	7.0	90.4	38.5	0.0	6.0	0.5	23.7
pl	20.9	100.0	20.8	0.0	11.0	3.5	26.0
bg	49.5	71.1	14.3	0.0	16.5	6.1	26.2
ka	0.0	95.0	34.0	12.0	13.0	12.0	27.7
nl	14.1	91.9	25.4	1.0	24.1	11.3	28.0
de	17.1	90.6	38.2	2.5	17.0	9.1	29.1
it	3.5	98.5	50.0	6.7	10.5	5.6	29.1
hu	6.1	98.5	33.5	14.0	28.0	10.6	31.8
hi	22.2	63.8	64.3	8.0	28.0	8.6	32.5
ru	32.0	95.0	33.8	7.1	24.0	3.0	32.5
fa	10.6	84.2	51.7	5.0	38.5	11.6	33.6
et	11.1	91.4	36.1	14.5	49.0	0.0	33.7
es	9.6	99.5	62.0	0.0	32.5	5.5	34.8
ms	4.0	97.9	84.0	5.5	18.5	0.5	35.1
ml	7.1	73.8	74.0	6.5	32.5	19.0	35.5
lt	12.7	84.2	67.6	25.2	22.5	2.5	35.8
el	6.1	98.0	38.3	18.1	45.0	16.5	37.0
af	47.0	78.7	35.5	1.0	56.0	21.2	39.9
ta	59.0	78.7	65.5	11.5	13.5	13.2	40.2
uk	16.2	94.9	39.8	26.9	51.5	14.1	40.6
pt	51.8	99.0	62.4	0.8	23.0	7.1	40.7
tl	0.0	100.0	66.5	12.5	58.5	13.6	41.8
id	6.5	98.0	77.0	0.0	42.0	33.5	42.8
ko	18.5	98.5	34.5	9.0	42.0	60.9	43.9
tr	100.0	84.4	72.5	0.5	21.0	14.1	48.8
pa	99.5	56.3	100.0	0.0	38.0	10.6	50.7
vi	13.6	99.0	80.0	10.0	100.0	2.6	50.9
te	35.2	92.0	68.0	28.0	36.0	55.0	52.4
ar	23.9	97.5	100.0	0.0	100.0	22.6	57.3
eu	100.0	98.5	66.0	17.5	25.5	38.7	57.7
bn	89.9	94.0	91.0	9.0	48.5	24.7	59.5
ur	90.9	57.5	76.1	16.5	100.0	18.7	59.9
my	99.0	86.0	83.0	9.5	93.5	0.0	61.8
kk	88.9	99.0	82.5	0.5	100.0	17.0	64.7
az	98.0	75.1	73.5	4.5	40.5	99.5	65.2
jv	9.0	100.0	82.0	3.5	100.0	100.0	65.8
mr	0.0	83.0	82.8	100.0	45.0	86.4	66.2
gu	100.0	100.0	100.0	39.0	90.0	31.8	76.8
ja	96.0	100.0	60.5	99.0	30.0	95.5	80.2
zh	94.4	100.0	33.0	100.0	94.5	71.2	82.2
sw	100.0	100.0	94.0	6.0	100.0	98.0	83.0
th	91.4	78.4	100.0	100.0	100.0	41.0	85.1
he	100.0	97.5	100.0	100.0	100.0	27.9	87.6
qu	91.9	100.0	100.0	98.0	95.5	97.0	97.1
ht	95.5	100.0	100.0	100.0	100.0	90.9	97.7
ha	100.0	100.0	99.5	100.0	100.0	91.5	98.5
yo	100.0	100.0	100.0	100.0	100.0	99.5	99.9
wo	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Avg	47.3	90.9	64.8	26.7	51.6	32.8	52.4

Table 5: Error rate of XLM-R fine-tuned on English SQuAD v1.1 on 6 CHECKLIST QA tests.

compared to the zero-shot setting (Hu et al., 2020). The new tasks are challenging for current models, which show relatively lower performance compared to other tasks. XCOPA presents a challenging classification task that requires cross-lingual common sense reasoning while the language-agnostic nature of Mewsli-X and LAReQA puts the

cross-lingual alignment of multilingual representations to the test. Analysis of the language-agnostic retrieval results show a large gap remains between cross-lingual and same-language test cases. XLM-R Large improves significantly over mBERT on the cross-lingual case in exchange for a slight drop for the same-language case. This points to XLM-R Large inducing more “strongly-aligned” representations (see Appendix I for details). The state-of-the-art mT5 improves performance on classification and QA tasks but performs less well on structured prediction and retrieval, highlighting settings where advances beyond scale are needed.⁶ Training on task-specific translations is beneficial in all cases and generally performs best, although improvements on QA tasks are diminishing. To obtain a more fine-grained understanding of the performance of current models, we conduct several analyses using our multilingual diagnostic suite.

5 Analyses

5.1 MULTICHECKLIST

We show the results of XLM-R fine-tuned on English SQuAD v1.1 on the 6 tests of MULTICHECKLIST in Table 5 (see Appendix F for the full results, example failure cases, and mBERT results). While mBERT’s average error rate is greater than 85% on 4/6 test categories, XLM-R demonstrates a substantially more robust cross-lingual understanding ability. XLM-R performs worst on tests in low-resource languages with limited or no pre-training data such as gu, ha, ht, qu, sw, wo, and yo and in languages with non-Latin scripts such as he, ja,

⁶Due to compute limitations, we extract mT5 embeddings by averaging the encoder outputs of a frozen mT5 model fine-tuned on SQuAD v1.1, as opposed to fine-tuning a dual encoder. For this reason, the mT5 language-agnostic retrieval scores are not directly comparable to those of mBERT and XLM-R.

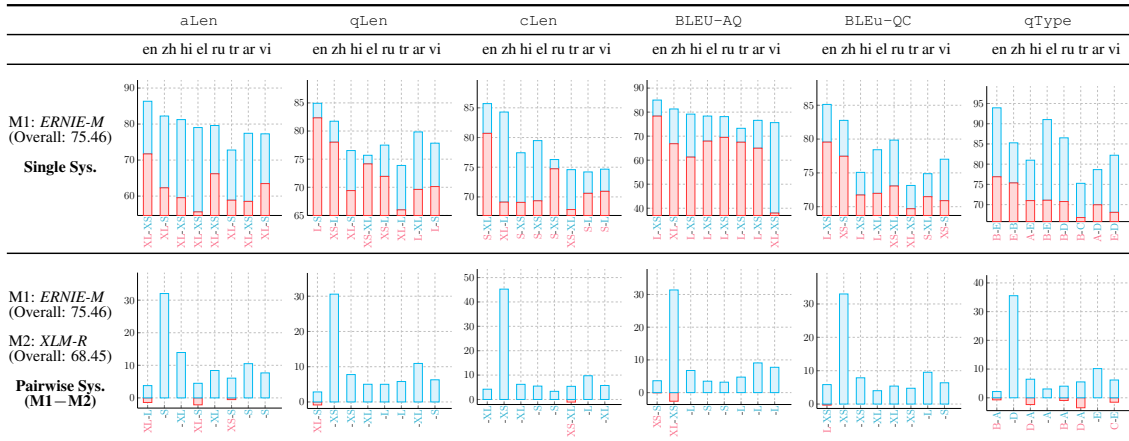


Table 6: Single and pairwise system diagnosis of ERNIE-M (M1) and XLM-R (M2) on XQuAD. “M1–M2” represents the performance difference between M1 and M2. We average “F1” and “Exact Match” of QA systems. We classify the attribute values into four *categories*: extra-small (XS), small (S), large (L) and extra-large (XL) values (see the Appendix for detailed interval information). The top 5 most frequent question types (qType) are: “what” (A), “how” (B), “who” (C), “when” (D) and “which” (E). In the *single system diagnosis* histogram, blue (red) x ticklabels represent the bucket category (e.g., XS) of a specific attribute on which a system achieved the best (worst) performance. In the *pairwise system diagnosis* histogram, blue (red) x ticklabels represent the bucket value of a specific attribute where system M1 surpasses (under-performs) M2 by the largest margin illustrated by a blue (red) bin. Blue-only x ticklabels (e.g., -D) indicate that M1 outperforms M2 in all categories of an attribute.

th, and zh. In addition, XLM-R displays interesting variation across languages, for instance failing in modeling comparisons in some languages, like Basque (eu), where it otherwise succeeds. We release the tests and test outputs to encourage deeper analysis and extension to other tasks and languages.

5.2 Nuanced Multilingual Evaluation

We showcase how nuanced multilingual evaluation enables us to perform *single* and *pairwise* system diagnosis on XQuAD in Table 6 (see Appendix K for analyses of the other tasks). We choose two systems: ERNIE-M, one of the top systems on XTREME, and XLM-R in eight languages: English, Chinese, Hindi, Greek, Russian, Turkish, Arabic, and Vietnamese (en, zh, hi, el, ru, tr, ar, vi).

Attributes We denote $(\mathbf{X}_c, \mathbf{X}_q, \mathbf{X}_a)$ as a tuple of a context, question and answer, and refer to cLen, qLen, aLen as their lengths (i.e., the number of tokens). We use BLEU (Papineni et al., 2002) to measure lexical overlap between $(\mathbf{X}_a, \mathbf{X}_q)$ and $(\mathbf{X}_q, \mathbf{X}_c)$ as BLEU-AQ and BLEU-QC. We report the top 5 most frequent question types (qType), which cover 85% of questions in the training set.

Single System Analysis For almost all languages, ERNIE-M achieves the highest performance on shorter answers (XS), but the worst performance on longer answers (XL). Especially in el, the performance difference between long and short an-

swers is more than 40 absolute points. The influence of question and context length is language-dependent. For example, in zh the system favors long questions and contexts while in hi, it is the opposite. If the answer is lexically similar to the question (larger BLEU-AQ), the system tends to make more mistakes in all eight languages. However, a higher lexical overlap between questions and contexts (BLEU-QC) is helpful for some languages: el, ru, ar. Surprisingly, ERNIE-M struggles to answer relatively frequent question types (i.e., what, and how), while it performs better on less frequent questions, indicating that although questions about person, place and choice are less frequent, they are easier than abstract questions.

Pairwise System Analysis Although ERNIE-M outperforms XLM-R by a large margin, it is surpassed by XLM-R on a few buckets. In en, XLM-R is better at dealing with longer answers and questions. In tr, XLM-R surpasses ERNIE-M on samples with shorter answers and contexts. In zh, XLM-R performs better when dealing with questions that are lexically similar to the answers.

6 Conclusions

Our analyses and experiments have shed light on important directions where scale alone is not sufficient such as “strong” alignment, syntactic transfer, fine-grained natural language understanding, and

answering of abstract questions. We encourage the development of better inductive biases, pre-training objectives, and evaluation resources. We make our data, translations, evaluation resources, and interactive leaderboard supporting detailed comparative analyses available to help the community gain a better understanding of multilingual models.

7 Ethical Considerations

7.1 Language representation

XTREME-R seeks to improve language representation and language diversity in NLP research, which has been identified as a large challenge (Joshi et al., 2020). We tried to cover a set of languages that is as diverse as possible, while still providing access to evaluation data in multiple tasks for each language. Despite this, XTREME-R has little representation of languages of the Americas and Africa due to a lack of labeled datasets for these languages. In addition, some languages included in XTREME-R with few data available online are only covered in a small number of datasets (see Table 7). To ameliorate this, we release training data of tasks translated into other languages, as well as the new MULTICHECKLIST. We reiterate the ongoing need for creating labeled datasets for diverse tasks in under-represented languages, to facilitate the development and evaluation of NLP models for such languages. We emphasize the importance of participatory research (✓ et al., 2020) as a modus operandi for such work in order to involve marginalized communities in the research process.

7.2 Leaderboard chasing

New benchmarks incentivize researchers to hill-climb on aggregate metrics (Ethayarajh and Jurafsky, 2020). In addition, new benchmarks create new opportunities for models to reach “superhuman” performance, which may lead people outside the field to erroneously conclude that some model has “solved language”. We hope that our inclusion of EXPLAINBOARD and MULTICHECKLIST help to prevent such a fallacy, by enabling more fine-grained evaluation that goes beyond a single aggregate metric.

7.3 Biases in multilingual models

Multilingual models have been shown to reflect biases similar to their monolingual counterparts (Zhao et al., 2020). In addition, multilingual models are biased towards languages with more pre-

training data (Hu et al., 2020). Zero-shot cross-lingual transfer additionally introduces a bias towards the source language (Søgaard et al., 2018; Anastasopoulos and Neubig, 2020). Due to the paucity of training data in other languages, we nevertheless focus on English-centric transfer and encourage future dataset creation efforts to include training data in multiple languages.

7.4 Environmental concerns

XTREME-R aims to enable efficient evaluation of multilingual models. To this end, we created a new dataset, Mewsl-X, that captures the essence of multilingual entity linking against a diverse knowledge base but is computationally cheaper to evaluate than the large-scale Mewsl-9 (Botha et al., 2020). Nevertheless, the models that perform best on benchmarks like XTREME-R are generally large-scale Transformer models pre-trained on large amounts of data, which comes at a high cost (Strubell et al., 2019). We thus particularly encourage the development of efficient methods to adapt existing models to new languages (Pfeifer et al., 2020) rather than training multilingual models entirely from scratch.

Acknowledgements

We thank Marco Tulio Ribeiro for advice on CHECKLIST. We are grateful to Laura Rimell and Jon Clark for valuable feedback on drafts of this paper, and to Dan Gillick for feedback on the Mewsl-X dataset design. We thank Hila Gonen, Bidisha Samantha, and Partha Talukdar for advice on Arabic, Bengali, and Hebrew CHECKLIST examples.

References

- Antonios Anastasopoulos and Graham Neubig. 2020. [Should all cross-lingual embeddings speak english?](#) In *Proceedings of ACL 2020*, pages 8658–8679.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020a. On the Cross-lingual Transferability of Monolingual Representations. In *Proceedings of ACL 2020*.
- Mikel Artetxe, Sebastian Ruder, Dani Yogatama, Gorka Labaka, and Eneko Agirre. 2020b. [A call for more rigor in unsupervised cross-lingual learning.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7375–7388, Online. Association for Computational Linguistics.

- Mikel Artetxe and Holger Schwenk. 2019. Massively Multilingual Sentence Embeddings for Zero-Shot Cross-Lingual Transfer and Beyond. *Transactions of the ACL 2019*.
- Peter Bakker, Aymeric Daval-Markussen, Mikael Parkvall, and Ingo Plag. 2011. Creoles are typologically distinct from non-creoles. *Journal of Pidgin and Creole Languages*, 26(1):5–42.
- Yonatan Belinkov and James Glass. 2019. Analysis Methods in Neural Language Processing: A Survey. *Transactions of the ACL*.
- Jan A. Botha, Zifei Shan, and Daniel Gillick. 2020. Entity Linking in 100 Languages. In *Proceedings of EMNLP 2020*, pages 7833–7845.
- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. 2021. [InfoXLM: An information-theoretic framework for cross-lingual language model pre-training](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3576–3588, Online. Association for Computational Linguistics.
- Hyung Won Chung, Thibault Févry, Henry Tsai, Melvin Johnson, and Sebastian Ruder. 2021. Rethinking Embedding Coupling in Pre-trained Language Models. In *Proceedings of ICLR 2021*.
- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. TyDi QA: A Benchmark for Information-Seeking Question Answering in Typologically Diverse Languages. In *Transactions of the Association of Computational Linguistics*.
- Jonathan H Clark, Dan Garrette, Iulia Turc, and John Wieting. 2021. [CANINE: Pre-training an Efficient Tokenization-Free Encoder for Language Representation](#). *arXiv preprint arXiv:2103.06874*.
- 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 ACL 2020*.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. XNLI: Evaluating cross-lingual sentence representations. In *Proceedings of EMNLP 2018*, pages 2475–2485.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of NAACL 2019*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Kawin Ethayarajh and Dan Jurafsky. 2020. [Utility is in the eye of the user: A critique of NLP leaderboards](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4846–4853, Online. Association for Computational Linguistics.
- Yuwei Fang, Shuohang Wang, Zhe Gan, Siqi Sun, and Jingjing Liu. 2021. [FILTER: An Enhanced Fusion Method for Cross-lingual Language Understanding](#). In *Proceedings of AAAI 2021*.
- ∀, Wilhelmina Nekoto, Vukosi Marivate, Tshinondiwa Matsila, Timi Fasubaa, Taiwo Fagbohunbe, Solomon Oluwole Akinola, Shamsuddeen Muhammad, Salomon Kabongo Kabenamualu, Salomey Osei, Freshia Sackey, Rubungo Andre Niyongabo, Ricky Macharm, Perez Ogayo, Orevaoghene Ahia, Musie Meressa Berhe, Mofetoluwa Adeyemi, Masabata Mokgesi-Seling, Lawrence Okegbemi, Laura Martinus, Kolawole Tajudeen, Kevin Degila, Kelechi Ogueji, Kathleen Siminyu, Julia Kreutzer, Jason Webster, Jamiil Toure Ali, Jade Abbott, Iro Orife, Ignatius Ezeani, Idris Abdulkadir Dangana, Herman Kamper, Hady Elsahar, Goodness Duru, Ghollah Kioko, Murhabazi Espoir, Elan van Biljon, Daniel Whitenack, Christopher Onyefuluchi, Chris Chinenye Emezue, Bonaventure F. P. Dossou, Blessing Sibanda, Blessing Bassey, Ayodele Olabiyi, Arshath Ramkilowan, Alp Öktem, Adewale Akinfaderin, and Abdallah Bashir. 2020. [Participatory research for low-resourced machine translation: A case study in African languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, Online.
- Jinlan Fu, Pengfei Liu, Graham Neubig, and Mo Tab. 2020. Interpretable Multi-dataset Evaluation for Named Entity Recognition. In *Proceedings of EMNLP 2020*, pages 6058–6069.
- Daniel Gillick, Alessandro Presta, and Gaurav Singh Tomar. 2018. End-to-end retrieval in continuous space. *arXiv preprint arXiv:1811.08008*.
- Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. Colorless green recurrent networks dream hierarchically. In *Proceedings of NAACL-HLT 2018*.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [DeBERTa: Decoding-enhanced BERT with Disentangled Attention](#). In *Proceedings of ICLR 2021*.
- Michael A. Hedderich, David Adelani, Dawei Zhu, Jesujoba Alabi, Udia Markus, and Dietrich Klakow. 2020. Transfer Learning and Distant Supervision for Multilingual Transformer Models: A Study on African Languages. In *Proceedings of EMNLP 2020*.
- 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 Generalization. In *Proceedings of ICML 2020*.
- Zhengbao Jiang, Antonios Anastasopoulos, Jun Araki, Haibo Ding, and Graham Neubig. 2020. **X-FACTR: Multilingual factual knowledge retrieval from pre-trained language models**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5943–5959, Online. Association for Computational Linguistics.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. The State and Fate of Linguistic Diversity and Inclusion in the NLP World. In *Proceedings of ACL 2020*.
- Divyanshu Kakwani, Anoop Kunchukuttan, Satish Golla, Gokul N.C., Avik Bhattacharyya, Mitesh M. Khapra, and Pratyush Kumar. 2020. **IndicNLPsuite: Monolingual corpora, evaluation benchmarks and pre-trained multilingual language models for Indian languages**. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4948–4961, Online. Association for Computational Linguistics.
- Phillip Keung, Yichao Lu, Julian Salazar, and Vikas Bhardwaj. 2020. **Don’t use English dev: On the zero-shot cross-lingual evaluation of contextual embeddings**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 549–554, Online. Association for Computational Linguistics.
- Simran Khanuja, Sandipan Dandapat, and Anirudh Srinivasan. 2020. **GLUECoS : An Evaluation Benchmark for Code-Switched NLP**. In *Proceedings of ACL 2020*, pages 3575–3585.
- Daniel Khashabi, Arman Cohan, Siamak Shakeri, Pedram Hosseini, Pouya Pezeshkpour, Malihe Alikhani, Moin Aminnaseri, Marzieh Bitaab, Faeze Brahman, Sarik Ghazarian, Mozhddeh Gheini, Arman Kabiri, Rabeeh Karimi Mahabadi, Omid Memarrast, Ahmadreza Mosallanezhad, Erfan Noury, Shahab Raji, Mohammad Sadegh Rasooli, Sepideh Sadeghi, Erfan Sadeqi Azer, Niloofar Safi Samghabadi, Mahsa Shafaei, Saber Sheybani, Ali Tazarv, and Yadollah Yaghoobzadeh. 2020. **ParsiNLU: A Suite of Language Understanding Challenges for Persian**. *arXiv preprint arXiv:2012.06154*.
- Anne Lauscher, Vinit Ravishankar, Ivan Vulić, and Goran Glavaš. 2020. **From Zero to Hero: On the Limitations of Zero-Shot Cross-Lingual Transfer with Multilingual Transformers**. In *Proceedings of EMNLP 2020*.
- Patrick Lewis, Barlas Oguz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2020. **MLQA: Evaluating cross-lingual extractive question answering**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7315–7330, Online. Association for Computational Linguistics.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenge Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. **XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018, Online. Association for Computational Linguistics.
- Tal Linzen. 2020. **How can we accelerate progress towards human-like linguistic generalization?** In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5210–5217, Online. Association for Computational Linguistics.
- Patrick Littell, David R Mortensen, Ke Lin, Katherine Kairis, Carlisle Turner, and Lori Levin. 2017. **Uriel and lang2vec: Representing languages as typological, geographical, and phylogenetic vectors**. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 8–14.
- Pengfei Liu, Jinlan Fu, Yang Xiao, Weizhe Yuan, Shuaichen Chang, Junqi Dai, Yixin Liu, Zihuiwen Ye, and Graham Neubig. 2021. **ExplainaBoard: An explainable leaderboard for NLP**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 280–289, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. **Roberta: A robustly optimized bert pretraining approach**. *arXiv preprint arXiv:1907.11692*.
- Fuli Luo, Wei Wang, Jiahao Liu, Yijia Liu, Bin Bi, Songfang Huang, Fei Huang, and Luo Si. 2020. **Veco: Variable encoder-decoder pre-training for cross-lingual understanding and generation**. *arXiv preprint arXiv:2010.16046*.
- Joakim Nivre, Mitchell Abrams, Željko Agić, Lars Ahrenberg, Lene Antonsen, Maria Jesus Aranzabe, Gashaw Arutie, Masayuki Asahara, Luma Ateyah, Mohammed Attia, et al. 2018. **Universal dependencies 2.2**.
- Xuan Ouyang, Shuohuan Wang, Chao Pang, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 2020. **ERNIE-M: Enhanced Multilingual Representation by Aligning Cross-lingual Semantics with Monolingual Corpora**. *arXiv preprint arXiv:2012.15674*.

- 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 ACL 2017*, pages 1946–1958.
- 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.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020. MAD-X: An Adapter-based Framework for Multi-task Cross-lingual Transfer. In *Proceedings of EMNLP 2020*.
- Jason Phang, Iacer Calixto, Phu Mon Htut, Yada Pruksachatkun, Haokun Liu, Clara Vania, Katharina Kann, and Samuel R. Bowman. 2020. [English intermediate-task training improves zero-shot cross-lingual transfer too](#). 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 557–575, Suzhou, China. Association for Computational Linguistics.
- Edoardo Maria Ponti, Goran Glavaš, Olga Majewska, Qianchu Liu, Ivan Vulić, and Anna Korhonen. 2020. [XCOPA: A multilingual dataset for causal commonsense reasoning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2362–2376, Online. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research*, 21:1–67.
- Afshin Rahimi, Yuan Li, and Trevor Cohn. 2019. Massively Multilingual Transfer for NER. In *Proceedings of ACL 2019*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. In *Proceedings of EMNLP 2016*.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. Beyond Accuracy: Behavioral Testing of NLP Models with CheckList. In *Proceedings of ACL 2020*, pages 4902–4912.
- Melissa Roemmele, Cosmin Adrian Bejan, and Andrew S Gordon. 2011. Choice of plausible alternatives: An evaluation of commonsense causal reasoning. In *AAAI spring symposium: logical formalizations of commonsense reasoning*, pages 90–95.
- Uma Roy, Noah Constant, Rami Al-Rfou, Aditya Barua, Aaron Phillips, and Yinfei Yang. 2020. [LARQA: Language-agnostic answer retrieval from a multilingual pool](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5919–5930, Online. Association for Computational Linguistics.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, Yejin Choi, Artificial Intelligence, and Computer Science. 2019. Social IQa: Commonsense Reasoning about Social Interactions Maarten Sap. In *Proceedings of EMNLP 2019*, pages 4453–4463.
- Anders Søgaard, Sebastian Ruder, and Ivan Vulić. 2018. [On the Limitations of Unsupervised Bilingual Dictionary Induction](#). In *Proceedings of ACL 2018*.
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. [Energy and Policy Considerations for Deep Learning in NLP](#). In *Proceedings of ACL 2019*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *Proceedings of NIPS 2017*.
- Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.
- Alex Wang, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2019a. SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems. In *Proceedings of NeurIPS 2019*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2019b. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In *Proceedings of ICLR 2019*.
- Xiangpeng Wei, Yue Hu, Rongxiang Weng, Luxi Xing, Heng Yu, and Weihua Luo. 2021. On Learning Universal Representations Across Languages. In *Proceedings of ICLR 2021*.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *Proceedings of NAACL-HLT 2018*.
- Bryan Willie, Karissa Vincentio, Genta Indra Winata, Samuel Cahyawijaya, Xiaohong Li, Zhi Yuan Lim, Sidik Soleman, Rahmad Mahendra, Pascale Fung, Syafri Bahar, and Ayu Purwarianti. 2020. IndoNLU: Benchmark and Resources for Evaluating Indonesian Natural Language Understanding. In *Proceedings of AACL-IJCNLP 2020*.
- Thomas Wolf, L Debut, V Sanh, J Chaumond, C Delangue, A Moi, P Cistac, T Rault, R Louf, M Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

Jieyu Zhao, Subhabrata (Subho) Mukherjee, Saghar Hosseini, Kai-Wei Chang, Ahmed Hassan Awadallah, Ahmed Hassan, Rdc Associates, Jieyu Zhao, Subhabrata (Subho) Mukherjee, Saghar Hosseini, Ahmed Awadallah, M Najafabadipour, J M Tuñas, A Rodríguez-González, and E Menasalvas. 2020. [Gender Bias in Multilingual Embeddings and Cross-lingual Transfer](#). In *Proceedings of ACL 2020*.

Appendix

A Details of XTREME models

All submissions to the XTREME leaderboard are large-scale Transformers (Vaswani et al., 2017) trained with masked language modeling (MLM), many of which extend monolingual models. Multilingual BERT (mBERT), XLM-RoBERTa (XLM-R; Conneau et al., 2020) and multilingual T5 (mT5; Xue et al., 2021) extend BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), and T5 (Raffel et al., 2020) respectively. Rebalanced mBERT (RemBERT; Chung et al., 2021) is a more efficient, scaled-up reparameterization of mBERT. These models have been pre-trained on unlabeled data in around 100 languages from Wikipedia (mBERT) and CommonCrawl. XLM-R was the strongest baseline in XTREME (Hu et al., 2020) and is the foundation for some subsequent work. It has been fine-tuned on English data of a related task prior to task-specific fine-tuning (STILTs; Phang et al., 2020). The following models furthermore propose new methods to leverage parallel data during pre-training or fine-tuning. FILTER (Fang et al., 2021), based on XLM-R, fuses representations in different languages. VECO (Luo et al., 2020) is a 24-layer encoder-decoder model that uses additional MLM variants during pre-training. T-URLv2 and HiCTL (Wei et al., 2021), based on InfoXLM (Chi et al., 2021) and XLM-R respectively, employ contrastive losses. ERNIE-M (Ouyang et al., 2020) incorporates back-translation into language modeling.⁷

B Task scores on XTREME

We show the performance of the models on the XTREME leaderboard broken down by language family on the remaining XTREME tasks in Figure 2.

C XTREME tasks retained in XTREME-R

XNLI The Cross-lingual Natural Language Inference corpus (Conneau et al., 2018) requires a model to determine whether a premise sentence entails, contradicts, or is neutral with respect to a hypothesis sentence. We use the crowd-sourced English data that was professionally translated to 14 other languages for evaluation and the MultiNLI (Williams et al., 2018) train set for training.

⁷We are not aware of the technical details of Polyglot and the anonymous submission.

UD-POS We employ the part-of-speech tagging data from the Universal Dependencies v2.7 treebanks (Nivre et al., 2018) covering 104 languages. We use the English training data for training and evaluate on the test sets of the target languages.

WikiANN-NER For named entity recognition, we use the WikiANN dataset (Pan et al., 2017), in which proper noun spans in Wikipedia text have been automatically annotated as either *location*, *person*, or *organization*. We use the balanced train, dev, and test splits from Rahimi et al. (2019).

XQuAD The Cross-lingual Question Answering Dataset (Artetxe et al., 2020a) requires identifying the answer to a question as a span in the corresponding paragraph. A subset of the English SQuAD v1.1 (Rajpurkar et al., 2016) dev set was professionally translated into ten other languages for XQuAD.

MLQA Similarly to XQuAD, the Multilingual Question Answering dataset (Lewis et al., 2020) is another cross-lingual question answering task. The evaluation data in seven languages was automatically mined from Wikipedia, annotations were crowd-sourced, and answer spans aligned. For both XQuAD and MLQA, we use their respective data for evaluation and train on SQuAD v1.1.

TyDiQA-GoldP We use the gold passage (GoldP) version of TyDiQA (Clark et al., 2020), a benchmark for information-seeking question answering, which covers nine typologically diverse languages. The GoldP version is a simplification of the primary task, using only the gold passage as context and excluding unanswerable questions. We use the English training data for training and evaluate on the test sets of the target languages.


Tatoeba We evaluate on the Tatoeba dataset (Artetxe and Schwenk, 2019), which consists of up to 1,000 English-aligned sentence pairs covering 122 languages. We find the nearest neighbor using cosine similarity. To make the setting more realistic, we move away from zero-shot retrieval and fine-tune models on SQuAD v1.1.

D Languages


Language characteristics We show a detailed overview of languages in XTREME-R including interesting typological differences in Table 7. Wikipedia information is taken from Wikipedia⁸ and linguistic information from WALS Online⁹.

⁸https://meta.wikimedia.org/wiki/List_of_Wikipedias


⁹<https://wals.info/languoid>



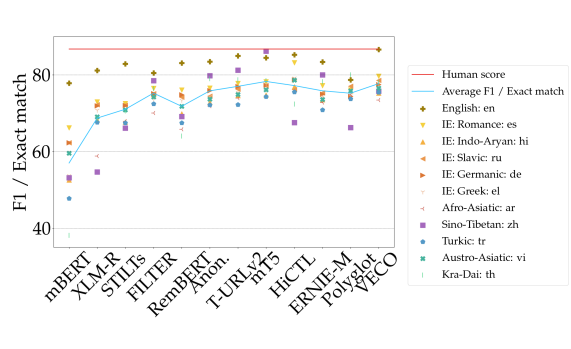
(a) Performance on XNLI




(b) Performance on PAWS-X




(c) Performance on WikiANN-NER



(d) Performance on XQuAD



(e) Performance on TyDiQA-GoldP



(f) Performance on BUCC

Figure 2: Performance of all models on the XTREME leaderboard on (a) XNLI, (b) PAWS-X, (c) WikiANN-NER, (d) XQuAD, (e) TyDiQA-GoldP, and (f) BUCC across language families.

XTREME-R includes members of the Afro-Asiatic, Austro-Asiatic, Austronesian, Dravidian, Indo-European, Japonic, Kartvelian, Kra-Dai, Niger-Congo, Sino-Tibetan, Turkic, Uralic, Creole¹⁰, and Quechuan language families as well as of two isolates, Basque and Korean.

Language diversity indices We measure the language diversity of XTREME-R according to the typology and language family indices of Ponti et al. (2020), which we show in Table 8 for XTREME-R, XTREME (Hu et al., 2020), and XGLUE (Liang et al., 2020). The typology index is based on the mean entropy of the distribution over 103 typological

¹⁰For simplicity, we treat Creole as a distinct language family (Bakker et al., 2011).

features from URIEL (Littell et al., 2017) across the languages while the family index consists of the number of distinct language families divided by the total number of languages. XTREME-R is similarly diverse while covering a larger number of languages.

E Mewslis-X Dataset

Mewslis-X is constructed specifically for XTREME-R and is a more carefully sampled variant of the Mewslis-9 dataset (Botha et al., 2020), derived from WikiNews in the same way. Compared to Mewslis-9, Serbian is dropped and Polish, Romanian and Ukrainian are added to obtain 11 languages, while the entity descriptions

Language	ISO 639-1 code	# Wikipedia articles (in millions)	Script	Language family	Diacritics / special characters	Extensive compounding	Bound words / clitics	Inflection	Derivation	# datasets with language
Afrikaans	af	0.09	Latin	IE: Germanic		✓				3
Arabic	ar	1.02	Arabic	Afro-Asiatic	✓		✓	✓		9
Azerbaijani*	az	0.18	Latin	Turkic					✓	2
Bulgarian	bg	0.26	Cyrillic	IE: Slavic	✓		✓	✓		4
Bengali	bn	0.08	Brahmic	IE: Indo-Aryan	✓	✓	✓	✓	✓	3
German	de	2.37	Latin	IE: Germanic		✓		✓		8
Greek	el	0.17	Greek	IE: Greek	✓	✓		✓		6
English	en	5.98	Latin	IE: Germanic						9
Spanish	es	1.56	Latin	IE: Romance	✓		✓			8
Estonian	et	0.2	Latin	Uralic	✓	✓		✓	✓	4
Basque	eu	0.34	Latin	Basque	✓		✓	✓	✓	3
Persian	fa	0.7	Perso-Arabic	IE: Iranian		✓				3
Finnish	fi	0.47	Latin	Uralic				✓	✓	4
French	fr	2.16	Latin	IE: Romance	✓		✓			4
Gujarati*	gu	0.03	Brahmic	IE: Indo-Aryan			✓			1
Hebrew	he	0.25	Jewish	Afro-Asiatic				✓		3
Hindi	hi	0.13	Devanagari	IE: Indo-Aryan	✓	✓	✓	✓	✓	7
Haitian Creole*	ht	0.06	Latin	Creole						1
Hungarian	hu	0.46	Latin	Uralic	✓	✓		✓	✓	3
Indonesian	id	0.51	Latin	Austronesian			✓	✓	✓	5
Italian	it	1.57	Latin	IE: Romance	✓		✓			4
Japanese	ja	1.18	Ideograms	Japonic			✓	✓		4
Javanese	jv	0.06	Brahmic	Austronesian	✓		✓			1
Georgian	ka	0.13	Georgian	Kartvelian				✓	✓	2
Kazakh	kk	0.23	Arabic	Turkic	✓			✓	✓	2
Korean	ko	0.47	Hangul	Koreanic		✓		✓	✓	4
Lithuanian*	lt	0.2	Latin	IE: Baltic	✓			✓		3
Malayalam	ml	0.07	Brahmic	Dravidian	✓	✓	✓	✓		2
Marathi	mr	0.06	Devanagari	IE: Indo-Aryan			✓	✓		3
Malay	ms	0.33	Latin	Austronesian			✓	✓		2
Burmese	my	0.05	Brahmic	Sino-Tibetan	✓	✓				1
Dutch	nl	1.99	Latin	IE: Germanic		✓				3
Punjabi*	pa	0.04	Brahmic	IE: Indo-Aryan	✓			✓		1
Polish*	pl	1.44	Latin	IE: Slavic	✓			✓		4
Portuguese	pt	1.02	Latin	IE: Romance	✓		✓			3
Cusco Quechua*	qu	0.02	Latin	Quechuan					✓	2
Romanian*	ro	0.42	Latin	IE: Romance	✓		✓	✓		4
Russian	ru	1.58	Cyrillic	IE: Slavic				✓		7
Swahili	sw	0.05	Latin	Niger-Congo			✓	✓	✓	4
Tamil	ta	0.12	Brahmic	Dravidian	✓	✓	✓	✓	✓	5
Telugu	te	0.07	Brahmic	Dravidian	✓	✓	✓	✓	✓	4
Thai	th	0.13	Brahmic	Kra-Dai	✓					7
Tagalog	tl	0.08	Brahmic	Austronesian	✓		✓	✓		2
Turkish	tr	0.34	Latin	Turkic	✓	✓		✓	✓	7
Ukrainian*	uk	1.06	Cyrillic	IE: Slavic				✓		4
Urdu	ur	0.15	Perso-Arabic	IE: Indo-Aryan	✓	✓	✓	✓	✓	4
Vietnamese	vi	1.24	Latin	Austro-Asiatic	✓					8
Wolof*	wo	0.002	Latin	Niger-Congo	✓					1
Yoruba	yo	0.03	Arabic	Niger-Congo	✓					2
Mandarin	zh	1.09	Chinese ideograms	Sino-Tibetan		✓				8

Table 7: Statistics about languages in XTREME-R. Languages belong to 14 language families and two isolates, with Indo-European (IE) having the most members. * indicates newly added languages. Diacritics / special characters: Language adds diacritics (additional symbols to letters). Compounding: Language makes extensive use of word compounds. Bound words / clitics: Function words attach to other words. Inflection: Words are inflected to represent grammatical meaning (e.g. case marking). Derivation: A single token can represent entire phrases or sentences.

	Range	XTREME-R	XTREME	XGLUE
Typology	[0, 1]	0.42	0.43	0.35
Family	[0, 1]	0.34	0.38	0.37

Table 8: Diversity indices with regard to typology and language family of XTREME-R, XTREME, and XGLUE.

to be retrieved range over all 50 languages in XTREME-R (Table 9). To broaden accessibility, the mention queries, candidates and Wikipedia-based training instances are all downsampled compared to the previous work.

Mention Extraction Viable mentions were taken as hyperlinks in the WikiNews articles pointing to Wikipedia entity pages (in any language) that could be mapped successfully to items in our base WikiData entity vocabulary. V_{base} is defined as the set of entities that have a Wikipedia page in any of the 50 languages in XTREME-R ($|V_{\text{base}}| \approx 12\text{M}$). The latter condition ensures that entity descriptions are available. We also extended the WikiData filter used by Botha et al. (2020) to additionally exclude entities that are instances of Wikipedia List Pages (Q13406463) or Wikimedia (Human) Disambiguation Pages (Q22808320, Q4167410), which are not of large interest for entity linking.

The resolved, viable mentions were filtered to drop duplicate surface mentions of an entity in the same article, and mentions of years (e.g. 2014), which are commonly linked in WikiNews but not of great interest. We then performed stratified sampling by both mention language and entity frequency bins, seeking uniform sizes across strata. Entity frequency is estimated as the number of times an entity is referenced on pages in the 50-language Wikipedia collection, and then binned into five intervals: $[0, 1)$, $[1, 10)$, $[10, 100)$, $[100, 1000)$, $[1000, \infty)$.

The resulting set of 15,000 test mentions covers 9,647 distinct gold entities (V_{gold}).

Candidate Set Doing vector encoding and nearest neighbor search over a full knowledge base of millions of entities is relatively time-consuming, but searching only among V_{gold} is also unrealistic. We thus strike a balance by defining a candidate set $V_{\text{cand}} \supset V_{\text{gold}}$, by sampling additional items from $V_{\text{base}} \setminus V_{\text{gold}}$, this time only stratified by entity frequency, to obtain $|V_{\text{cand}}| = 1,000,000$.

Each entity $e \in V_{\text{cand}}$ is assigned a single description: we randomly sample a language from among the $L_e \leq 50$ Wikipedia pages corresponding to e , and take the page’s first sentence as entity description.

Fine-tuning Data The fine-tuning data constitutes 115K English-only (mention, entity)-pairs, which were sampled from English Wikipedia hyperlinks that map to $V_{\text{base}} \setminus V_{\text{cand}}$. Sampling is according to the natural distribution. The Mewsli-X evaluation

setting in XTREME-R is thus doubly zero-shot: no candidate or test entities are observed during fine-tuning, nor is non-English text.

F MULTICHECKLIST

General statistics Creating MULTICHECKLIST involved translating around 550 words (templates and fill-in values) into 49 languages. Some languages required a larger number of words due to the creation of additional templates (Russian required translating 1043 words). In total, the translation effort cost \$4,360. For each test category and each language, we automatically generate 200 test cases. Depending on the number of possible variations for each test, each test case can consist of 2 (Comparisons) to 12 (Intensifiers) examples.¹¹

Instructions to translators We sent the following guidelines to annotators for the translation:

1. We would like you to translate the following templates and their corresponding fill-in values into other languages. Each template contains some text enclosed in curly brackets $\{ \}$. These are the names of the fields that will be substituted in the template. We ask you not to translate the text within such curly brackets.
2. Have a look at the text in the other fields to get a better sense what can be substituted in the template. For instance, by referring to the lines with “adj”, we can see that “{first_name} is {adj[0]} than {first_name1}” is a comparison between two people.
3. If there is not a literal translation or the same translation has already been used for another word, feel free to use a translation that is similar in meaning.
4. If the translation of the template differs based on the substituted words, please create multiple translations for the template and indicate which substituted words correspond to it. For instance, if one translation of the template assumes that some substituted words have male gender but others have female gender, create a separate translation of the template that conforms to the female gender.
5. In one template, {p.p1} and {p.p2} refer to a property (either “shape”, “color”, “size”,

¹¹Note that the number of variations does not correlate with the test success rate, e.g. the “Job vs nationality” test has 8 variations for each test case and the highest success rate overall.

“age”, or “material”). {p.v1} and {p.v2} refer to an attribute of each property such as “old”, “new”, “red”, “blue”, etc.

Challenges of template translation We highlight some of the challenges and linguistic phenomena we encountered during the process of creating MULTICHECKLIST in the following. Unless specified otherwise, we create separate templates to disambiguate each phenomenon.

- **Gender agreement:** Adjectives and nationalities need to be declined to match the gender of their referring expression. To keep the translation effort manageable and avoid creating separate templates for each gender, we control for gender and restrict fill-in values to male names for the affected tests (3/6). We sample genders equally for the other tests. We welcome dedicated test suites analyzing multilingual gender bias as future extensions.
- **Declination:** In Russian, animal and vehicle names require Accusative and Nominative in different cases.
- **Normalization:** For appropriate substitution, fill-in values often need to include articles. We normalize answers and predictions by removing language-specific articles in order to ensure a consistent comparison.
- **Names:** Our use of names based on data in Wikidata leads to certain biases. Names that are more common in Wikidata are more likely to be chosen. In some cases, names in Wikidata are not written in the native script. Japanese names from Wikidata are often written in hiragana or katakana rather than kanji. Our choice of using the first name is also not applicable to all languages. In Japanese, people are usually not referred to by their first name, e.g. Masanori Suzuki would be called Suzuki-san instead of Masanori.
- **Declension of names:** In some languages, a suffix is appended to a name depending on its spelling. For instance, in Turkish the suffix changes based on the vowel of the last syllable, e.g. Ahmet'in, Ali'nin, Umut'un, Şeyma'nın, Özge'nin, etc., and Ahmet'ten, Ali'den, Umut'tan, Şeyma'dan, Özge'den, etc. In Finnish, names are appended with a variation of “lla”, e.g. Peterillä, Lisalla, Mattilla, etc.

- **Professions:** Words for certain professions are gendered (e.g. *waiter/waitress*), so they only occur with male or female names.
- **Question syntax:** In some languages, the syntax of the question changes depending on the property or adjective one asks about.
- **Syntax of adjectives:** In some languages, the syntax changes depending on what adjective is used. In German, the translations of “happy”, “excited”, and “passionate” require different prepositions.
- **Stilted language:** Some text appears stilted when values are filled into the translated templates. For instance, the question “*どちらの方が冷静でないですか。*” is an unusual way to do negation in Japanese; if directly translated to English, it would mean “Who is more not calm?”.

We tried to address most of these challenges by instructing translators to create additional templates to disambiguate linguistic phenomena and by consolidating different templates programmatically. However, as this process was relatively labor-intensive, we recommend the usage of morphologically aware templates similar to Jiang et al. (2020) for future work. Note that morphologically aware templates may not be able to resolve some of the finer linguistic differences. For this reason, we also advise working closely with native speakers to design tests that reflect natural language as closely as possible.

Full results We show the full results of XLM-R and mBERT on the MULTICHECKLIST tests in Tables 10 and 11 respectively. mBERT only shows limited amounts of cross-lingual taxonomic knowledge. While it is able to distinguish between job and nationality and animals and vehicles in some languages, it fails to do this consistently across all languages. In addition, it completely fails to distinguish between different properties and intensifiers and is not able to perform comparisons. In contrast, while XLM-R struggles with intensifiers, it demonstrates the other capabilities much more consistently across languages.

Example failure cases We provide example failure cases of XLM-R on a subset of languages in Table 12. We will publicly release a comprehensive list of failure cases for XLM-R and mBERT, the complete tests and model outputs for further analysis.

G Hyper-parameters

mBERT We use the cased version, which covers 104 languages, has 12 layers, 768 hidden units per layer, 12 attention heads, a 110k shared WordPiece vocabulary, and 110M parameters.¹² The model was trained using Wikipedia data in all 104 languages, oversampling low-resource languages with an exponential smoothing factor of 0.7. We generally fine-tune mBERT for two epochs, with a training batch size of 32 and a learning rate of $2e-5$. We build on the Transformers library (Wolf et al., 2019) for training on each task.

XLM-R We use the XLM-R Large version that covers 100 languages, uses a 200k shared BPE vocabulary, and has been trained with masked language modelling.¹³ We fine-tune XLM-R generally for two epochs with a learning rate of $3e-5$ and an effective batch size of 16. We use the Transformers library for training XLM-R on all tasks.

mT5 We use the publicly released mT5-XXL version that has nearly 13 billion parameters with a vocabulary size 250k (Xue et al., 2021). It has been trained on multilingual C4 (mC4) corpus which has 6.3 trillion tokens spanning 101 languages¹⁴. For all downstream tasks, we fine-tune mT5-XXL for 10k steps with a constant learning rate of 0.001, dropout rate of 0.1 and a batch size of 2^{17} tokens. For early stopping, we save checkpoints every 200 steps and choose the checkpoint with the highest performance on the validation set.

H Metadata

We intend to ask each submission to XTREME-R for relevant metadata. Such metadata includes the number of parameters, amount of pre-training data, amount of fine-tuning data, etc. We are doing this to enhance transparency and to increase utility of our benchmark for practitioners with varying needs. As a first step in this direction, we provide information about the number of parameters and the amount of monolingual and parallel pre-training data used by all submissions to XTREME in Table 13. Note that the different systems report their training data in different ways (e.g. number of tokens, number of examples, size of the data). We plan to standardize this by asking submissions to

¹²<https://github.com/google-research/bert/blob/master/multilingual.md>

¹³<https://github.com/facebookresearch/XLM>

¹⁴<https://www.tensorflow.org/datasets/catalog/c4#c4multilingual>

XTREME-R to report training data in terms of number of tokens seen.

I Language-agnostic Retrieval Results

The multiway cross-language nature of Mewsli-X and LAReQA enables closer analysis of model performance by input and target *language pairs*. Mewsli-X can directly be split by language pair as it has a single correct target per input mention. For LAReQA, we follow the “Limit to One Target” strategy of Roy et al. (2020): instead of asking the model to retrieve all correct answers in one pass, we evaluate on each target separately, with all the other correct answers removed from the candidate pool, allowing us to report splits by language pair.

Table 14 summarizes these pairwise mAP@20 scores (here, micro-averaged), showing that XLM-R Large improves substantially over mBERT on the cross-lingual case (+38% on Mewsli-X and +137% for LAReQA) in exchange for a slight drop for the same-language case. Even so, performance on the cross-lingual case is still low at 29–36 mAP@20, and remains a challenging area for future work. Figures 3 and 4 show the detailed breakdowns.

J Detailed results

We show the detailed results for each task and language in Tables 15 (XNLI), 16 (XCOPA), 17 (UD-POS), 18 (WikiANN-NER), 19 (XQuAD), 20 (MLQA), 21 (TyDiQA-GoldP), 22 (Tatoeba), 23 (Mewsli-X), and 24 (LAReQA).

K Nuanced Multilingual Evaluation

We perform nuanced multilingual evaluations by categorizing testing examples into different attribute buckets and measuring the system performance on each attribute bucket. In the following, we describe the available attributes for tasks in XTREME-R and provide additional analysis on different attributes.

K.1 Attribute Definition

QA We denote $(\mathbf{X}_c, \mathbf{X}_q, \mathbf{X}_a)$ as a tuple of the corresponding context, question and answer, and refer to $cLen$, $qLen$, $aLen$ as their lengths (i.e., the number of tokens). We use BLEU (Papineni et al., 2002) to measure the lexical overlap between $(\mathbf{X}_a, \mathbf{X}_q)$ and $(\mathbf{X}_q, \mathbf{X}_c)$ as BLEU-AQ and BLEU-QC. We classify questions based on their first tokens and report the top 5 most frequent question types as $qType$ (i.e., what, how, when, where,

	ar	de	en	es	fa	ja	pl	ro	ta	tr	uk
ar	78.8	12.5	6.5	0.7	35.8	0.0	2.3	3.1	2.5	4.9	9.1
de	0.0	89.2	88.4	61.1		22.9	62.3	65.9	0.0	66.1	9.8
en	4.4	86.7	83.9	73.2	0.0	8.7	60.7	72.2	0.0	73.2	8.5
es	10.7	72.2	70.5	88.5		15.1	64.0	44.4	0.9	49.1	28.5
fa	39.9	4.1	4.5	6.0	89.2	5.9	0.2	7.2	0.2	1.2	6.6
ja	1.7	11.6	14.9	13.5	0.0	88.7	11.6		0.4	0.0	1.4
pl	0.6	58.3	64.9	65.4	0.0	4.9	85.8	49.3	0.2	74.5	7.6
ro	8.8	63.1	70.5	69.5			55.3	77.1	1.0	32.9	19.4
ta	10.0		0.3	0.0		10.0		0.0	76.5		
tr	6.7	55.9	58.0	66.3	3.1	1.8	57.9	21.2	0.0	82.0	30.6
uk	2.3	14.5	20.7	15.3		4.5	25.0	22.7	0.8	11.7	86.2
af										62.9	
az	0.0	75.8				0.0	43.3	77.4	0.7	71.9	
bg	6.9		12.5				38.3		0.0		54.4
bn	0.0		0.7					0.0	0.3	0.0	
el	0.0		5.9	0.5		5.9	1.5	1.1	3.1	6.5	5.4
et	0.0	60.2	51.0				48.5	26.5	0.0	70.0	5.8
eu	3.3	60.5		72.7				74.0	0.0	18.5	17.9
fi	3.7	61.5	76.4	63.7		10.3	66.7	18.3	0.0	47.5	9.7
fr	3.2	67.6	69.7	70.2	0.0	12.0	51.0	58.8	2.0	55.0	5.6
gu					0.0						
he	4.0	0.1	19.2	9.7	0.8	17.0	14.1	0.8	1.5	0.0	9.0
hi	1.2									1.0	
ht										0.0	
hu	4.2	63.9		54.1		3.0		65.3	0.6	53.7	14.5
id	3.2	65.4	58.8	46.6	6.8	1.9	67.4	36.0	0.4	65.1	16.0
it	9.0	77.6	75.3	79.0	4.5	1.2	62.6	51.6	5.9	59.5	10.3
jv	0.0							95.5	0.0		
ka	0.5								1.0	25.0	11.0
kk						5.3	8.4	0.0		66.5	
ko	1.3	5.9	14.0	11.0	2.0	13.5	2.4	0.6	1.4	5.1	5.8
lt	6.0	44.8		40.8		24.0	42.0	26.2	2.8	17.7	16.4
ml	0.0							0.0	1.4	0.0	0.0
mr	8.8		1.1	0.0	9.6			0.0	1.2	11.1	
ms	1.8	78.8	55.4	44.0		5.5	25.3	63.4	1.6	72.0	25.5
my				0.0				0.0	0.0		
nl	4.7	68.2	63.2	57.7		0.0	56.9	1.6	0.0	64.7	12.3
pa					35.2				10.0		
pt	1.7	66.1	57.0	86.5	2.0	14.8	71.8	78.8	0.0	49.1	11.8
qu									33.9	0.0	
ru	1.7	24.4	22.8	33.7	6.9	0.3	14.7	20.6	0.8	17.4	66.4
sw										0.0	
te										0.0	
th			0.0	0.0		0.0	8.3			0.0	
tl						9.9		15.9	0.0		
ur	54.2	2.0						0.3	0.0	1.6	
vi	9.1	67.2	54.0	60.9	32.7	20.6	55.2	63.1	1.7	51.0	13.8
wo								0.0			
yo								23.6	0.0		
zh	3.6	14.0	17.4	17.2	1.6	66.8	8.6	26.6	0.0	4.7	0.6

(a) mBERT

	ar	de	en	es	fa	ja	pl	ro	ta	tr	uk
ar	78.3	27.1	21.1	24.1	40.0	1.9	15.9	4.9	7.3	46.8	24.8
de	12.8	86.6	84.8	70.7		23.9	66.6	58.6	5.6	79.2	33.8
en	16.9	86.8	79.6	71.1	6.0	12.0	63.7	82.2	10.2	84.7	31.9
es	32.4	68.7	78.6	81.7		11.5	56.3	46.5	15.2	59.1	43.5
fa	73.0	15.9	28.3	12.0	92.6	4.8	16.5	6.7	8.4	15.2	8.9
ja	8.9	31.9	40.0	32.3	5.4	86.6	11.3		21.1	10.8	16.9
pl	16.8	75.0	67.3	71.2	11.0	17.6	85.9	50.8	2.2	73.5	26.6
ro	22.0	70.3	56.3	64.8			55.2	72.8	4.9	63.6	38.0
ta	19.4		4.4	20.9		4.9		0.5	86.1		
tr	25.8	65.2	60.9	64.9	5.0	8.7	63.7	36.7	8.3	80.2	38.3
uk	21.8	39.5	58.7	35.2		15.5	39.1	46.1	30.7	43.7	86.9
af										66.7	
az	20.0	76.9				10.8	39.5	68.7	3.9	63.8	
bg	22.4		54.5				47.3		23.5		76.9
bn	36.8		12.5					13.3	2.7	17.9	
el	13.0		26.8	8.5		14.6	50.0	10.2	11.9	31.2	18.4
et	0.0	49.0	46.6				58.8	32.6	10.7	51.8	39.0
eu	8.2	60.2		78.1				67.0	12.6	18.7	38.7
fi	28.7	65.7	92.2	61.9		10.1	71.7	74.4	3.3	57.1	47.6
fr	17.1	70.7	74.2	69.9	20.8	10.2	62.6	66.1	25.1	71.6	38.0
gu						1.3				20.0	
he	9.1	4.7	34.2	22.1	14.2	19.8	24.9	2.0	6.2	4.3	28.4
hi	6.0									4.0	
ht										0.0	
hu	28.1	74.3		58.3		17.9		63.0	0.0	80.4	36.1
id	15.6	65.4	55.2	49.6	30.6	6.0	69.9	28.0	12.9	69.1	48.4
it	22.8	80.3	86.2	79.1	28.1	16.4	57.4	78.6	8.2	76.7	27.3
jv	1.4								34.8	17.5	
ka	25.6									15.4	37.1
kk								22.3	19.7	6.2	78.6
ko	14.1	22.8	16.6	37.9	14.6	15.3	11.6	1.0	7.1	9.6	26.7
lt	10.8	56.2		53.3		32.0	62.3	76.0	18.0	59.7	56.1
ml	10.1							8.7	21.3	1.8	8.6
mr	21.5		16.1	15.5	25.4				14.7	19.8	20.8
ms	11.4	80.3	57.7	44.4		16.3	24.5	60.7	21.1	68.9	41.3
my				3.6				2.5	7.3		
nl	14.9	71.0	67.2	67.2		1.5	71.9	2.1	18.8	61.5	32.5
pa						32.9				13.1	
pt	22.0	67.8	54.0	87.7	33.0	8.8	72.9	81.7	18.2	74.2	37.5
qu									51.1	0.0	
ru	25.5	46.3	47.9	43.7	13.7	16.6	40.2	22.0	22.7	46.1	75.9
sw										15.4	
te										27.0	
th			21.5	21.8		3.2	33.7			13.8	
tl						14.2		13.6	16.4		
ur	74.2	25.0						3.3	38.3	5.0	
vi	7.6	77.9	54.0	55.4	11.1	25.0	44.0	30.9	4.8	34.5	36.0
wo								0.0			
yo								10.1	1.0		
zh	14.8	20.0	26.1	28.8	12.3	51.0	12.2	7.7	5.7	10.7	14.1

(b) XLM-R Large

Figure 3: Mewslis-X results broken down by mention language (columns) and entity description language (rows), only showing combinations with at least 10 test items.




Figure 4: LARQA results broken down by question language (rows) and answer language (columns) for “Limit to One Target” inference.

which), which cover 85% of questions in the training set. We list the six attributes as follows.

- $\phi_{aLen}(\mathbf{X}_a) = \text{len}(\mathbf{X}_a)$,
- $\phi_{qLen}(\mathbf{X}_q) = \text{len}(\mathbf{X}_q)$,
- $\phi_{cLen}(\mathbf{X}_c) = \text{len}(\mathbf{X}_c)$,
- $\phi_{BLEU_{aq}}(\mathbf{X}_a, \mathbf{X}_q) = \text{BLEU}(\mathbf{X}_a, \mathbf{X}_q)$,
- $\phi_{BLEU_{qc}}(\mathbf{X}_q, \mathbf{X}_c) = \text{BLEU}(\mathbf{X}_q, \mathbf{X}_c)$,
- $\phi_{type}(\mathbf{X}_q) = \text{type}(\mathbf{X}_q)$, *question type*.

Structured Prediction Given a sentence \mathbf{X} , we define the i -th word token as x_i and a span of words in the range of $[i, j)$ as $\mathbf{X}_{i:j}$ in the sentence. We then define five attributes including the label of a span (tag), the token length of a sentence (sLen), the token length of an entity span (eLen), the character length of an entity span (tLen) and the relative token position of an entity (rPos) in the sentence as follows.

- $\phi_{\text{tag}}(\mathbf{X}_{i:j}) = \text{label}(\mathbf{X}_{i:j})$
- $\phi_{\text{sLen}}(\mathbf{X}) = \text{len}(\mathbf{X})$
- $\phi_{\text{eLen}}(\mathbf{X}_{i:j}) = \text{len}(\mathbf{X}_{i:j})$
- $\phi_{\text{tLen}}(\mathbf{X}_{i:j}) = |\mathbf{X}_{i:j}|$
- $\phi_{\text{rPos}}(\mathbf{X}_{i:j}) = i/\phi_{\text{sLen}}(\mathbf{X})$, *relative position*

where $|x|$ represents the number of characters.

K.2 Attribute Buckets

We bucket all test examples into different attribute buckets for a given attribute. Specifically, for an attribute defined for a task, we measure the attribute value of the test examples (see Section K.1), then determine N attribute buckets for all test examples ($N = 4$ by default), and finally we measure the system performance on the test examples falling in each attribute interval to observe the performance change over different attribute buckets. Since the attribute values can be either continuous (e.g., answer length aLen) or discrete (e.g., question type qType), we perform different strategies for creating attribute buckets for them.

Continuous Attribute Values For the attributes with continuous attribute values (e.g., ϕ_{aLen} , ϕ_{qLen} , ϕ_{cLen} , $\phi_{\text{BLEU}_{aq}}$, $\phi_{\text{BLEU}_{qc}}$), we divide the test examples into different intervals where the numbers of the test samples in all attribute intervals are equal.

Discrete Attribute Values For the attribute with discrete attribute values (e.g., ϕ_{type}), test samples with the same type are put into the same attribute bucket.

Table 25 and 26 show the detailed attribute intervals for each category on the XQuAD task and WikiANN-NER task, respectively.

K.3 Additional Nuanced Analysis

K.3.1 WikiANN-NER

Table 27 and 28 illustrate the single system diagnosis of ERNIE-M and XLM-R respectively on the WikiANN-NER task in three languages (i.e., `en`, `es`, `fr`). We make the following observations.

ERNIE-M In Table 27, first, we observe that the effects of some attributes for ERNIE-M are language-independent. For example, based on the attribute `rPos`, the system is good at predicting entities located within the first 1/3 part of the English sentences, while it is relatively bad at predicting entities within the first 1/3 part of the sentences for other languages. Second, the system favors long sentences based on the attribute `sLen`. We even observe that performance increases as the sentence length increases on `es` and `fr`. Third, across all languages, the system performs relatively bad at predicting long entities (`eLen`) and entities belonging to the organization class (`tag`). Finally, the system is good at predicting sentences with fewer entities based on the attribute for entity density (`eDen`).

XLM-R In Table 28, we observe that the influence of some attributes such as `sLen`, `eLen`, `eDen` with respect to the system performance are similar between ERNIE-M and XLM-R, although ERNIE-M performs significantly better than XLM-R at generalizing its predictions on `es`, `fr`.

K.3.2 QA

Table 29 shows the pairwise system analysis of ERNIE-M and T-URLv2 for the XQuAD task. We find that although the overall performance of T-URLv2 outperforms ERNIE-M, it is surpassed by ERNIE-M on a few buckets. For example, in `zh`, ERNIE-M is better at dealing with samples that have long answers, long questions, and a high lexical overlap between questions and answers. In `ru`, ERNIE-M is better at dealing with samples with long answers, long questions, and lower lexical overlap between questions and answers, questions and contexts.

K.4 EXPLAINABOARD Demonstration

Figure 5 shows the interface of EXPLAINABOARD containing possible selection options to observe the fine-grained analysis for submitted systems on XTREME. We also demonstrate how to perform *Single System* and *Pair Systems* analysis on Figure 6 and 7 respectively.

Specifically, to generate a fine-grained overview, we first select models in the table, then click one of the three *Analysis Buttons*, which generates a fine-grained analysis such as in Figure 6 (*single system analysis*) and Figure 7 (*pair-wise system analysis*).

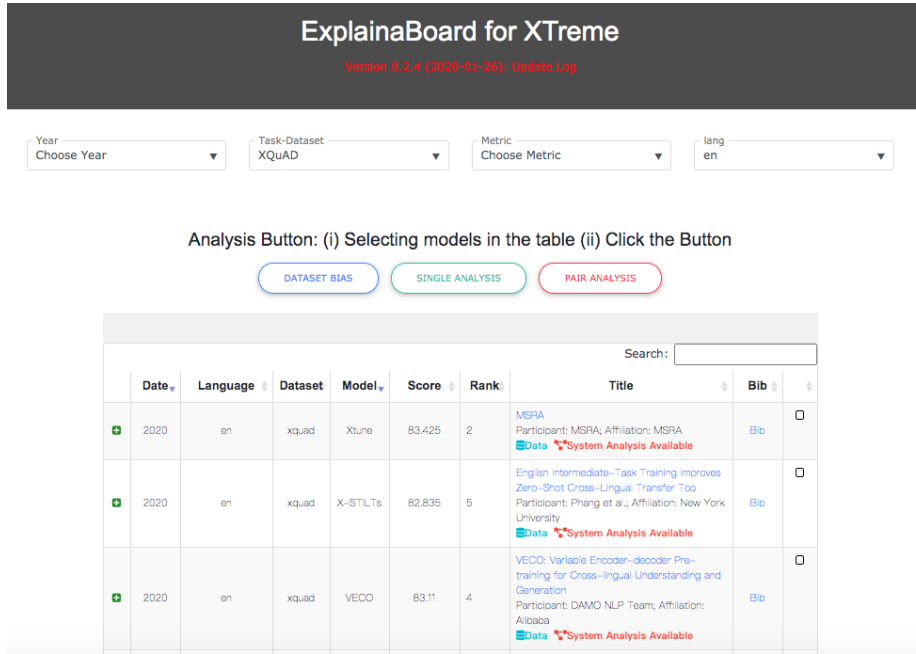


Figure 5: Overall user interface. Four top-down lists at the top are used to filter the entries in the table by the *publication year, task, metric* and *languages*. Three analysis buttons (e.g., *DATASET BIAS, SINGLE ANALYSIS, PAIRWISE ANALYSIS*) are used to perform three different fine-grained evaluations. Each row in the table represents the performance of a system on a specific dataset and a specific language. Relevant pieces of information such as the paper title also are provided.



Figure 6: Single system analysis. Each histogram represents the fine-grained evaluation results of a given system, which are broken down based on a pre-defined attribute (e.g., answer length).




Figure 7: Pairwise system analysis. Each histogram illustrates the fine-grained performance gap between system 1 and system 2, which has been broken down by different pre-defined attributes (e.g., answer length).

	ar	de	en	es	fa	ja	pl	ro	ta	tr	uk	Total
ar	185	23	64	36	12	16	17	28	42	69	24	516
de	40	485	57	56	7	34	81	53	34	81	38	966
en	88	107	470	162	18	86	124	31	62	171	108	1,427
es	61	42	43	410	7	35	41	26	26	24	35	750
fa	51	38	25	27	32	21	27	25	45	30	26	347
ja	22	34	30	24	12	586	32	6	24	14	26	810
pl	59	37	43	32	10	27	552	68	32	34	81	975
ro	28	24	19	21	6	8	20	107	34	19	30	316
ta	10	4	19	11	2	10	7	129	62	6	6	266
tr	60	17	16	21	13	14	21	13	15	347	17	554
uk	44	23	22	29	9	11	53	38	31	23	381	664
af	5	5	5	6	1	3	2	3	9	15	7	61
az	10	18	5	8	5	10	10	22	12	39	8	147
bg	20	9	11	7	5	9	15	9	17	5	32	139
bn	48	8	26	8	3	7	9	14	43	13	6	185
el	14	9	17	11	4	45	11	22	33	27	21	214
et	15	18	13	9	9	8	11	20	14	12	25	154
eu	15	28	8	23	6	8	8	35	32	13	14	190
fi	29	22	35	30	8	13	23	23	58	30	21	292
fr	50	59	69	65	12	46	60	15	29	56	50	511
gu	8	0	3	1	10	3	0	8	15	2	1	51
he	28	48	23	18	18	34	17	22	37	34	28	307
hi	12	4	4	7	7	3	6	0	21	6	6	76
ht	3	5	5	1	2	6	3	3	11	1	2	42
hu	12	21	9	11	2	12	9	23	27	19	22	167
id	22	24	24	13	10	11	11	17	35	29	42	238
it	47	48	43	53	20	16	33	21	17	31	40	369
jv	10	2	3	2	2	4	5	11	24	5	7	75
ka	24	4	8	7	6	4	5	9	13	25	13	118
kk	5	8	9	7	6	3	14	29	15	9	11	116
ko	28	33	27	39	19	62	20	52	69	36	20	405
lt	24	19	9	10	8	10	12	60	18	27	11	208
ml	14	7	7	8	5	2	7	14	33	11	11	119
mr	20	5	19	10	11	5	8	22	46	12	4	162
ms	19	13	13	13	4	47	10	27	35	14	11	206
my	7	5	8	14	3	2	4	10	24	7	6	90
nl	33	45	33	43	7	17	33	224	38	17	25	515
pa	3	0	4	3	2	12	1	2	10	2	3	42
pt	50	23	42	59	23	29	49	13	22	21	18	349
qu	6	3	1	1	9	0	3	22	11	7	1	64
ru	53	32	44	34	15	21	39	11	15	40	147	451
sw	5	6	5	5	9	9	1	8	13	7	7	75
te	1	3	3	4	3	4	6	6	27	6	2	65
th	5	4	16	15	2	28	12	1	12	6	7	108
tl	8	6	1	6	6	11	3	17	31	5	5	99
ur	14	10	4	8	7	6	6	43	27	14	5	144
vi	60	17	12	18	11	10	14	22	105	22	18	309
wo	1	2	0	1	1	6	3	15	5	1	2	37
yo	9	4	6	3	5	4	6	16	38	5	2	98
zh	55	39	38	36	14	71	17	67	21	25	28	411
Total	1,440	1,450	1,420	1,446	428	1,449	1,481	1,482	1,469	1,474	1,461	15,000

Table 9: Composition of Mewsl-X test set by mention language (columns) and entity description language (rows).

Lang.	Comparisons	Intensifiers	Properties	Job vs Nationality	Animal vs Vehicles	Animal vs Vehicles 2
af	47.0	78.7	35.5	1.0	56.0	21.2
ar	23.9	97.5	100.0	0.0	100.0	22.6
az	98.0	75.1	73.5	4.5	40.5	99.5
bg	49.5	71.1	14.3	0.0	16.5	6.1
bn	89.9	94.0	91.0	9.0	48.5	24.7
de	17.1	90.6	38.2	2.5	17.0	9.1
el	6.1	98.0	38.3	18.1	45.0	16.5
en	7.0	90.4	38.5	0.0	6.0	0.5
es	9.6	99.5	62.0	0.0	32.5	5.5
et	11.1	91.4	36.1	14.5	49.0	0.0
eu	100.0	98.5	66.0	17.5	25.5	38.7
fa	10.6	84.2	51.7	5.0	38.5	11.6
fi	10.7	79.8	34.1	4.0	9.2	1.5
fr	4.0	98.0	15.4	0.8	10.0	13.5
gu	100.0	100.0	100.0	39.0	90.0	31.8
ha	100.0	100.0	99.5	100.0	100.0	91.5
he	100.0	97.5	100.0	100.0	100.0	27.9
hi	22.2	63.8	64.3	8.0	28.0	8.6
ht	95.5	100.0	100.0	100.0	100.0	90.9
hu	6.1	98.5	33.5	14.0	28.0	10.6
id	6.5	98.0	77.0	0.0	42.0	33.5
it	3.5	98.5	50.0	6.7	10.5	5.6
ja	96.0	100.0	60.5	99.0	30.0	95.5
jv	9.0	100.0	82.0	3.5	100.0	100.0
ka	0.0	95.0	34.0	12.0	13.0	12.0
kk	88.9	99.0	82.5	0.5	100.0	17.0
ko	18.5	98.5	34.5	9.0	42.0	60.9
lt	12.7	84.2	67.6	25.2	22.5	2.5
ml	7.1	73.8	74.0	6.5	32.5	19.0
mr	0.0	83.0	82.8	100.0	45.0	86.4
ms	4.0	97.9	84.0	5.5	18.5	0.5
my	99.0	86.0	83.0	9.5	93.5	0.0
nl	14.1	91.9	25.4	1.0	24.1	11.3
pa	99.5	56.3	100.0	0.0	38.0	10.6
pl	20.9	100.0	20.8	0.0	11.0	3.5
pt	51.8	99.0	62.4	0.8	23.0	7.1
qu	91.9	100.0	100.0	98.0	95.5	97.0
ru	32.0	95.0	33.8	7.1	24.0	3.0
sw	100.0	100.0	94.0	6.0	100.0	98.0
ta	59.0	78.7	65.5	11.5	13.5	13.2
te	35.2	92.0	68.0	28.0	36.0	55.0
th	91.4	78.4	100.0	100.0	100.0	41.0
tl	0.0	100.0	66.5	12.5	58.5	13.6
tr	100.0	84.4	72.5	0.5	21.0	14.1
uk	16.2	94.9	39.8	26.9	51.5	14.1
ur	90.9	57.5	76.1	16.5	100.0	18.7
vi	13.6	99.0	80.0	10.0	100.0	2.6
wo	100.0	100.0	100.0	100.0	100.0	100.0
yo	100.0	100.0	100.0	100.0	100.0	99.5
zh	94.4	100.0	33.0	100.0	94.5	71.2
Avg	47.3	90.9	64.8	26.7	51.6	32.8

Table 10: Error rate of XLM-R fine-tuned on English SQuAD v1.1 on 6 CHECKLIST QA tests across all 50 languages.

Lang.	Comparisons	Intensifiers	Properties	Job vs Nationality	Animal vs Vehicles	Animal vs Vehicles 2
af	92.5	100.0	91.5	30.0	86.5	52.0
ar	91.9	100.0	100.0	16.5	100.0	79.4
az	100.0	100.0	94.0	6.5	95.5	98.5
bg	99.0	100.0	92.9	0.0	80.0	74.5
bn	100.0	100.0	100.0	100.0	100.0	88.4
de	97.5	100.0	90.2	10.0	100.0	16.2
el	95.9	100.0	95.5	61.3	99.0	83.0
en	100.0	100.0	96.5	0.0	25.0	15.1
es	100.0	100.0	87.8	0.0	28.0	10.1
et	78.9	100.0	97.6	76.5	100.0	3.0
eu	100.0	100.0	98.5	91.5	79.0	77.4
fa	95.5	100.0	99.5	45.5	99.5	69.8
fi	90.8	100.0	74.6	15.5	89.7	42.6
fr	100.0	100.0	83.6	19.3	17.0	15.5
gu	100.0	100.0	100.0	92.0	100.0	99.0
ha	91.0	100.0	100.0	100.0	100.0	94.0
he	100.0	100.0	100.0	100.0	100.0	75.6
hi	99.0	100.0	93.8	12.5	94.0	80.8
ht	72.0	100.0	100.0	100.0	100.0	68.2
hu	89.8	100.0	94.0	62.0	81.0	57.3
id	99.0	100.0	80.0	8.0	94.0	32.0
it	99.5	100.0	95.7	20.2	39.5	31.3
ja	100.0	100.0	99.0	100.0	100.0	78.4
jv	94.5	100.0	99.5	50.0	100.0	91.5
ka	100.0	100.0	99.5	58.5	47.5	53.0
kk	98.0	100.0	97.5	5.5	100.0	74.0
ko	97.5	100.0	75.0	14.5	99.5	55.3
lt	92.9	100.0	99.5	37.0	94.0	37.0
ml	98.0	100.0	92.0	100.0	100.0	97.0
mr	100.0	100.0	100.0	100.0	100.0	99.5
ms	100.0	100.0	99.5	8.0	100.0	31.3
my	100.0	100.0	93.5	84.5	100.0	52.8
nl	96.5	100.0	92.2	45.5	41.5	8.8
pa	100.0	100.0	100.0	54.0	100.0	100.0
pl	99.0	100.0	88.9	33.5	81.0	53.8
pt	98.0	100.0	97.6	2.5	63.5	26.8
qu	100.0	100.0	100.0	100.0	96.5	98.5
ru	96.0	100.0	93.0	73.1	79.0	17.0
sw	100.0	100.0	100.0	55.5	100.0	86.5
ta	90.0	100.0	100.0	21.0	100.0	95.4
te	88.4	100.0	95.5	91.0	100.0	98.5
th	99.5	100.0	100.0	100.0	100.0	100.0
tl	99.5	100.0	100.0	96.5	100.0	100.0
tr	100.0	100.0	100.0	70.5	77.0	67.7
uk	95.8	100.0	85.6	81.5	65.0	37.9
ur	100.0	100.0	100.0	99.5	100.0	100.0
vi	97.5	100.0	90.5	7.5	100.0	31.1
wo	59.5	100.0	100.0	100.0	100.0	91.5
yo	99.0	100.0	99.0	100.0	100.0	65.3
zh	100.0	100.0	100.0	100.0	100.0	92.4
Avg	95.8	100.0	95.3	55.1	87.0	64.1

Table 11: Error rate of mBERT fine-tuned on English SQuAD v1.1 on 6 CHECKLIST QA tests across all 50 languages.

Lang.	Comparisons	Intensifiers	Properties	Job vs Nationality	Animal vs Vehicles	Animal vs Vehicles 2
af	C: Don is kleiner as William. Q: Wie is kleiner? A: Don P: William.	C: Tim is optimisties oor die projek. John is nogal optimisties oor die projek. Q: Wie is die mins optimisties oor die projek? A: Tim P: John	C: Daar is 'n kers in die kamer. Die kers is blou en ysteragtig. Q: Watter materiaal is die kers? A: ysteragtig P: blou en ysteragtig.		C: Keith het 'n brandweerwa en 'n konyn. Q: Watter dier het Keith? A: 'n konyn P: brandweerwa en 'n konyn.	C: Steven het 'n akkedis gekoop. Carl het 'n pancelwa gekoop. Q: Wie het 'n voertuig gekoop? A: Carl P: Steven
ar	C: ياسين اضعف من سليم. Q: من اقل ضعيف? A: سليم P: ياسين	C: سمع للغاية ايجابي بشأن المشروع. ايجابي بشأن المشروع. Q: من الاقل ايجابي بشأن المشروع? A: انطوان P: سمع	C: يوجد مروحة في الغرفة المروحة هو ابيض. Q: اي مقاس هو المروحة? A: ضئيل P: أزرق وضئيل.		C: محمد لديه بطة وسيارة دفع رباعي. Q: ما السيارة الموجودة لدى محمد? A: سيارة دفع رباعي. P: دفع رباعي.	C: سعد اشترى افعى. اشترى دراجة بخارية. Q: من اشترى السيارة? A: Lotfi P: سعد
az	C: Oleg Narimandan daha tündür. Q: Kim daha az tündür? A: Narimandan P: Oleg Narimandan	C: Yuri layihə ilə bağlı olaraq haddən çox coşquludur. Vasif layihə ilə bağlı olaraq coşquludur. Q: Layihə ilə bağlı olaraq az kim coşquludur? A: Vasif P: Yuri	C: Otaqda saat var. Saatavidir və dəmir. Q: material saat necədir? A: dəmir P:avidir və dəmir.	C: Tofiq amerikalıdır. Tofiq inzibatçısıdır. Q: Tofiqin işi nədir? A: inzibatçısıdır P: amerikalıdır. Tofiq inzibatçısıdır.	C: Samadin dovşanı və yük maşını var. Q: Samadin hansı nəqliyyat vasitəsi var? A: yük maşını P: dovşanı və yük maşını	C: Vitali kartənkəşliyi aldı. Emil kartənkəşliyi aldı. Q: Nəqliyyat vasitəsini kim aldı? A: Vitali P: Emil kartənkəşliyi
bg	C: Боян е по-странен от Веселин. Q: Кой е по-малко странен? A: Веселин P: Боян	C: Ljubomir е донякъде ентузиазиран за проекта. Лазар е ентузиазиран за проекта. Q: Кой е най-малко ентузиазиран за проекта? A: Ljubomir P: Лазар	C: Има канале в стаята. Каналето е черно и стъклено. Q: Какъв материал е каналето? A: стъклено P: черно и стъклено.		C: Младен има хамстер и пожарен камيون. Q: Какво превозно средство има Младен? A: пожарен камيون P: хамстер	C: Любомир купи бик. Антон купи колело. Q: Кой купи превозно средство? A: Антон P: Любомир
bn	C: মুহাম্মাদ খালেদ এর চেয়ে বেশি জ্ঞানী। Q: কে কম জ্ঞানী? A: খালেদ P: মুহাম্মাদ খালেদ	C: সোহেল এই প্রকল্প সম্পর্কে সংশয়বাদী। আব্দুল এই প্রকল্প সম্পর্কে অভ্যস্ত সংশয়বাদী। Q: কে এই প্রকল্প সম্পর্কে সবচেয়ে কম সংশয়বাদী? A: সোহেল P: আব্দুল	C: রুমের মধ্যে একটি ল্যাম্প আছে। ল্যাম্পটি হলো বাদামি ও কাঠের। Q: কোন উপাদান হলো ল্যাম্পটি? A: কাঠের P: বাদামি ও কাঠের।	C: হালিম একজন একজন ওয়েটস এবং চীনা। Q: হালিম এর জাতীয়তা কী? A: চীনা P: ওয়েটস এবং চীনা।	C: টনি এর একটি সাপ ও একটি ভ্যান আছে। Q: টনি এর কোন গাড়ি আছে? A: একটি ভ্যান P: একটি সাপ ও একটি ভ্যান	C: আখতার একটি হ্যামস্টার কিনেছেন। আবদুল কাদের একটি কার কিনেছেন। Q: কে একটি প্রাণী কিনেছেন? A: আখতার P: আবদুল কাদের
de	C: Heinrich ist dunkler als Harald. Q: Wer ist weniger dunkel? A: Harald P: Heinrich ist dunkler als Harald.	C: Helmut ist etwas beunruhigt über das Projekt. Alfred ist höchst beunruhigt über das Projekt. Q: Wer ist am wenigsten beunruhigt über das Projekt? A: Helmut P: Alfred	C: Es gibt eine Kamera in dem Raum. Die Kamera ist violett und winzig. Q: Welche Farbe hat die Kamera? A: violett P: violett und winzig.	C: Hugo ist ein Vertreter und nigerianisch. Q: Was ist Hugos Beruf? A: ein Vertreter P: Vertreter und nigerianisch.	C: Joachim hat einen Leguan und ein SUV. Q: Was für ein Fahrzeug hat Joachim? A: ein SUV P: Leguan	C: Markus hat einen Leguan gekauft. Oliver hat einen Minivan gekauft. Q: Wer hat ein Fahrzeug gekauft? A: Oliver P: Markus
el	C: Ο Παύλος είναι πιο παράξενος από τον Ευάγγελος. Q: Ποιος είναι λιγότερο παράξενος; A: Ευάγγελος P: Παύλος	C: Ο Μάριος είναι ενθουσιασμένος για το έργο. Ο Αντώνης είναι υπερβολικά ενθουσιασμένος για το έργο. Q: Ποιος είναι λιγότερο ενθουσιασμένος για το έργο; A: Μάριος P: Ο Αντώνης	C: Υπάρχει ένας τοίχος στο δωμάτιο. Ο τοίχος είναι μεγάλη και τριγωνική. Q: Τι μέγεθος έχει ο τοίχος; A: μεγάλη P: μεγάλη και τριγωνική.	C: Ο Βασίλειος είναι ερευνητής και Μπαγκλαντεσιανός. Q: Τι δουλειά κάνει ο Βασίλειος; A: ερευνητής P: ερευνητής και Μπαγκλαντεσιανός.	C: Ο Πάρης έχει μια γάτα και ένα μίνιβαν. Q: Τι όχημα έχει ο Πάρης; A: ένα μίνιβαν P: Πάρης έχει μια γάτα και ένα μίνιβαν.	C: Ο Ελίας αγόρασε ένα μίνιβαν. Ο Κώστας αγόρασε έναν ταύρο. Q: Ποιος αγόρασε ένα όχημα; A: Ελίας P: Ο Ελίας αγόρασε ένα μίνιβαν. Ο Κώστας

Table 12: Example failure cases of XLM-R on a subset of languages. Each failure case consists of a context (C), a question (Q), an answer (A), and XLM-R's prediction (P).

Model	Number of parameters (in millions)	Pre-training data	
		Monolingual data	Parallel data
mBERT	178	85 GB	N/A
XLm-R (large)	559	6.3T tokens	N/A
MMTE	190	N/A	25B pairs
mT5	13,000	1T tokens	N/A
RemBERT	575	1.8T tokens	N/A
X-STILTS	559	6.3T tokens	N/A
FILTER	559	6.3T tokens	N/A
VECO	559	1.3 TB	6.4M pairs
T-URLv2 + StableTune	559	2.1 TB	42 GB
ERNIE-M	559	1.5 TB	69 GB

Table 13: Metadata for the current submissions to XTREME. Note that monolingual pre-training data is reported in either number of tokens or size of the data (in GB/TB). The amount of parallel data is reported in either number of pairs or size of the data (in GB/TB).

Subset \ Model	Mewslis-X		LArEQA	
	mBERT	XLm-R Large	mBERT	XLm-R Large
All	40.2	47.1	16.3	31.3
Same language	85.7	83.6	58.2	57.1
Different languages	25.8	35.5	12.1	28.7

Table 14: Language-agnostic retrieval results (mAP@20) broken down by whether the query and answer languages are the same or different.

Model	en	ar	bg	de	el	es	fr	hi	ru	sw	th	tr	ur	vi	zh	avg
mBERT	81.7	66.0	69.2	71.1	66.8	74.9	74.2	60.7	70.2	49.3	54.7	61.2	58.2	70.5	69.0	65.4
XLm-R	88.7	77.2	83.0	82.5	80.8	83.7	82.2	75.6	79.1	71.2	77.4	78.0	71.7	79.3	78.2	79.2
mT5	92.6	84.5	87.0	87.3	86.9	88.4	87.4	82.3	84.3	80.6	81.2	83.4	79.8	84.0	83.0	84.8
<i>mBERT translate-train</i>	80.8	73.6	76.6	77.4	75.7	78.1	77.4	71.9	75.2	69.4	70.9	75.3	67.2	75.0	74.1	74.6
<i>mBERT translate-train-all</i>	81.9	73.8	77.6	77.6	75.9	79.1	77.8	70.7	75.4	70.5	70.0	74.3	67.4	77.0	77.6	75.1

Table 15: XNLI results (accuracy) across languages.

Model	et	ht	id	it	qu	sw	ta	th	tr	vi	zh	Avg
mBERT	54.6	51.8	55.4	57.6	54.0	52.2	55.6	52.0	55.4	60.4	67.6	56.1
XLm-R	68.2	52.6	80.6	71.0	52.8	61.8	73.8	74.4	72.0	76.2	78.0	69.2
mT5	77.5	72.1	81.1	75.9	54.4	74.1	75.9	78.3	78.1	76.9	79.5	74.9
<i>mBERT translate-train</i>	57.4	55.6	60.6	63.4	47.8	51.6	54.8	56.4	58.8	59.6	64.0	57.3
<i>mBERT translate-train-all</i>	56.0	54.4	59.0	60.4	51.0	56.8	55.4	58.2	56.8	63.6	65.4	57.9

Table 16: XCOPA results (accuracy) across languages.

Model	af	ar	bg	de	el	en	es	et	eu	fa	fi	fr	he	hi	hu	id	it	ja	kk	ko
mBERT	86.1	54.0	85.2	85.6	80.4	95.4	85.8	79.5	59.3	66.2	78.0	83.4	55.3	67.1	78.3	80.6	88.2	49.7	69.8	49.2
XLm-R Large	89.7	68.1	88.6	88.5	86.5	96.1	89.1	87.2	74.3	73.9	86.3	88.4	68.2	74.4	83.4	82.9	89.8	29.4	79.3	54.0
	mr	nl	pt	ru	ta	te	th	tl	tr	ur	vi	yo	zh	lt	pl	uk	wo	ro	Avg	
mBERT	66.6	88.6	86.8	85.1	68.6	75.2	39.1	68.5	68.3	58.0	53.4	53.4	61.2	77.9	79.9	80.4	28.9	77.3	70.9	
XLm-R Large	85.3	89.5	89.3	89.7	77.3	85.3	47.7	75.0	75.4	67.2	56.8	22.8	40.8	84.5	84.8	85.8	27.6	85.5	75.0	

Table 17: UD-POS results (F1 score) across languages.

Model	ar	he	vi	id	jv	ms	tl	eu	ml	ta	te	af	nl	en	de	el	bn	hi	mr	ur	fa	fr	it	pt	es
mBERT	43.9	56.6	72.1	62.9	64.6	70.9	73.7	65.0	53.5	53.0	47.4	76.0	81.9	84.5	78.0	68.8	70.3	66.7	57.3	35.6	50.1	79.1	81.3	79.8	72.2
XLm-R large	43.7	54.1	77.2	52.3	58.7	69.8	72.2	62.1	65.8	56.9	52.3	77.7	84.3	84.6	78.0	77.2	76.3	71.0	64.2	54.1	61.1	79.1	81.1	79.6	68.8
	bg	ru	ja	ka	ko	th	sw	yo	my	zh	kk	tr	et	fi	hu	qu	pl	uk	az	lt	pa	gu	ro	Avg	
mBERT	77.2	63.5	28.4	65.0	59.1	2.3	72.7	49.4	49.1	43.3	49.0	72.1	77.0	77.4	75.0	57.4	79.9	69.3	65.6	74.3	37.0	47.3	72.1	62.7	
XLm-R large	81.2	71.5	18.3	68.9	58.0	1.5	69.9	41.8	50.9	25.8	49.9	78.9	78.0	78.6	79.3	50.4	81.3	73.1	69.2	76.9	46.8	60.7	79.6	64.4	

Table 18: WikiANN-NER results (F1 score) across languages.

Model	en	es	de	el	ru	tr	ar	vi	th	zh	hi	avg
mBERT	84.5	75.1	73.2	62.9	71.3	53.7	62.2	70.1	43.5	59.6	59.5	65.1
XLm-R Large	87.4	82.7	80.9	80.7	80.5	76.1	74.4	79.2	75.4	55.0	77.0	77.2
mT5	90.2	84.6	82.3	82.8	78.8	76.5	80.3	83.3	74.7	81.7	81.7	81.5
<i>mBERT translate-train</i>	83.5	80.2	75.6	70.0	75.0	68.9	68.0	75.6	36.9	66.2	69.6	70.0
<i>mBERT translate-train-all</i>	86.0	82.4	78.8	74.2	78.1	70.6	71.0	78.5	38.1	67.7	71.3	72.4

Table 19: XQuAD results (F1) across languages.

Model	en	es	de	ar	hi	vi	zh	avg
mBERT	80.7	66.2	60.2	51.6	49.9	60.2	60.3	61.3
XLm-R Large	83.9	74.4	70.3	67.0	70.8	74.2	68.4	72.7
mT5	86.4	76.2	73.1	70.2	75.3	76.5	71.4	75.6
<i>mBERT translate-train</i>	80.2	70.0	64.4	55.0	60.1	65.7	63.9	65.6
<i>mBERT translate-train-all</i>	80.7	71.3	66.0	58.9	62.4	67.9	66.0	67.6

Table 20: MLQA results (F1) across languages.

Model	en	ar	bn	fi	id	ko	ru	sw	te	avg
mBERT	69.4	61.7	53.5	57.4	63.2	57.6	56.5	59.7	46.2	58.4
XLm-R Large	69.2	66.1	59.1	64.7	73.8	58.0	62.2	66.4	59.1	64.3
mT5	83.1	82.4	83.6	81.2	84.5	73.2	78.7	87.2	83.6	81.9
<i>mBERT translate-train</i>	75.3	61.5	31.9	62.6	68.6	53.2	53.1	61.9	27.4	55.1
<i>mBERT translate-train-all</i>	73.2	71.8	49.7	68.1	72.3	58.6	64.3	66.8	53.3	64.2

Table 21: TyDiQA-GoldP results (F1) across different languages.

Model	ar	he	vi	id	jv	tl	eu	ml	ta	te	af	nl	de	el	bn	hi	mr	ur	fa	fr	it
mBERT	32.7	46.8	66.2	57.7	18.5	17.5	31.0	19.2	23.5	26.1	49.6	66.1	79.7	29.1	20.7	38.6	23.7	38.4	49.9	67.1	65.6
XLm-R	68.3	77.6	91.0	88.4	28.8	60.8	58.6	83.6	65.8	80.8	74.9	90.0	96.6	76.6	67.6	88.9	70.4	76.5	85.3	87.5	82.4
	pt	es	bg	ru	ja	ka	ko	th	sw	zh	kk	tr	et	fi	hu	az	lt	pl	uk	ro	Avg
mBERT	74.4	71.8	52.8	65.1	50.8	21.1	43.9	15.0	12.1	74.9	30.4	38.4	34.0	42.7	44.2	37.2	35.4	54.1	55.8	54.8	41.3
XLm-R	89.8	89.5	84.4	86.1	75.5	66.2	81.9	80.5	31.3	79.5	63.8	84.5	73.4	86.5	83.1	74.7	77.5	88.0	82.6	89.6	76.2

Table 22: Tatoeba results (accuracy) across different languages.

Model	ar	de	en	es	fa	ja	pl	ro	ta	tr	uk	avg
mBERT	15.3	61.0	54.8	59.4	13.5	44.2	57.7	27.5	4.1	49.9	37.0	38.6
XLm-R Large	28.7	64.8	59.7	62.0	24.6	45.0	62.1	30.8	14.9	59.4	51.2	45.7

Table 23: Mewsli-X results (mean average precision@20) across different input languages.

Model	ar	de	el	en	es	hi	ru	th	tr	vi	zh	avg
mBERT	17.0	29.3	16.3	31.3	30.8	12.3	27.2	5.8	17.9	25.3	24.2	21.6
XLm-R Large	34.6	42.8	38.8	46.2	43.7	38.2	42.5	39.5	41.3	40.9	39.8	40.7

Table 24: LAReQA results (mean average precision@20) across different question languages.

Bucket	aLen	qLen	cLen	BLEU_AQ	BLEU_QC	aLen	qLen	cLen	BLEU_AQ	BLEU_QC
	en					zh				
XS	[1, 2]	[3, 8]	[25, 91]	0	[0, 1.64e-08]	1	[1, 13]	[3, 8]	0	[0, .0075]
S	(2, 5]	(8, 11]	(91, 111]	(0, .016]	(1.64e-08, 2.8e-06]	(1, 6]	(13, 17]	(8, 11]	(0, .0118]	(.0075, .01]
L	(5, 13]	(11, 16]	(111, 156]	(.016, .128]	(2.8e-06, 7.92e-05]	(6, 15]	(17, 22]	(11, 17]	(.0118, .0542]	(.01, .013]
XL	(13, 25]	(16, 29]	(156, 509]	(.128, .146]	(7.92e-05, .0589]	(15, 72]	(22, 52]	(17, 144]	(.0542, .328]	(.013, .097]
hi					el					
XS	[1, 2]	[3, 8]	[28, 104]	0	[0, 2.02e-08]	[1, 2]	[3, 8]	[29, 97]	0	[4.99e-39, 9.58e-09]
S	(2, 5]	(8, 11]	(104, 131]	(0, .015]	(2.02e-08, 2.91e-06]	(2, 5]	(8, 11]	(97, 120]	(0, .014]	(9.58e-09, 1.32e-06]
L	(5, 13]	(11, 15]	(131, 171]	(.015, .19]	(2.91e-06, 7.29e-05]	(5, 11]	(11, 16]	(120, 162]	(.014, .019]	(1.32e-06, 4.37e-05]
XL	(13, 29]	(15, 53]	(171, 575]	(.019, .24]	(7.29e-05, .0628]	(11, 26]	(16, 30]	(162, 554]	(.019, .12]	(4.37e-05, .033]
ru					tr					
XS	[1, 2]	[3, 6]	[25, 86]	0	[1.34e-39, 3.24e-09]	[1, 2]	[2, 6]	[18, 75]	0	[0, 6.37e-09]
S	(2, 5]	(6, 8]	(86, 104]	(0, .017]	(3.24e-09, 6.2e-07]	(2, 5]	(6, 8]	(75, 94]	(0, .016]	(6.37e-09, 9.33e-07]
L	(5, 12]	(8, 11]	(104, 140]	(.017, .023]	(6.2e-07, 2.07e-05]	(5, 11]	(8, 11]	(94, 130]	(.016, .057]	(9.33e-07, 2.67e-05]
XL	(12, 23]	(11, 26]	(140, 470]	(.023, .287]	(2.07e-05, .0294]	(11, 20]	(11, 20]	(130, 414]	(.057, .143]	(2.67e-05, .0408]
ar					vi					
XS	[1, 2]	[2, 7]	[24, 83]	0	[0, 4.72e-09]	[1, 2]	[3, 10]	[39, 127]	0	[0, 1.47e-08]
S	(2, 4]	(7, 9]	(83, 104]	(0, .0243]	(4.72e-09, 8.51e-07]	(2, 4]	(10, 13]	(127, 160]	(0, .0139]	(1.47e-08, 2.43e-06]
L	(4, 9]	(9, 12]	(104, 138]	(.0243, .293]	(8.51e-07, 3.39e-05]	(4, 8]	(13, 17]	(160, 213]	(.0139, .135]	(2.43e-06, 7.76e-05]
XL	(9, 22]	(12, 38]	(138, 432]	(.293, .246]	(3.39e-05, .0266]	(8, 34]	(17, 40]	(213, 763]	(.135, .189]	(7.76e-05, .0546]

Table 25: Detailed attribute intervals for each category on the XQuAD task.

Bucket	en				es				fr			
	sLen	eLen	rPos	eDen	sLen	eLen	rPos	eDen	sLen	eLen	rPos	eDen
XS	[3,5]	1	0	[0,0.33]	[3,5]	1	0	[0,0.33]	[1,5]	1	0	[0.027,0.33]
S	(5,8]	2	(0,0.33]	(0.33,0.56]	(5,7]	2	(0,0.25]	(0.33,0.67]	(5,7]	2	(0,0.33]	(0.33,0.64]
L	(8,16]	(2,4]	(0.33,0.62]	(0.56,1]	(7,18]	(2,4]	(0.25,0.6]	(0.67,0.86]	(7,16]	(2,4]	(0.33,0.67]	(0.64,1]
XL	(16,1e+03]	(4,1e+03]	(0.62,1]	-	(18,1e+03]	(4,1e+03]	(0.6,1]	(0.86,1]	(16,1e+03]	(4,1e+03]	(0.67,1]	-

Table 26: Detailed attribute intervals for each category on WikiANN-NER task in en, es, and fr.

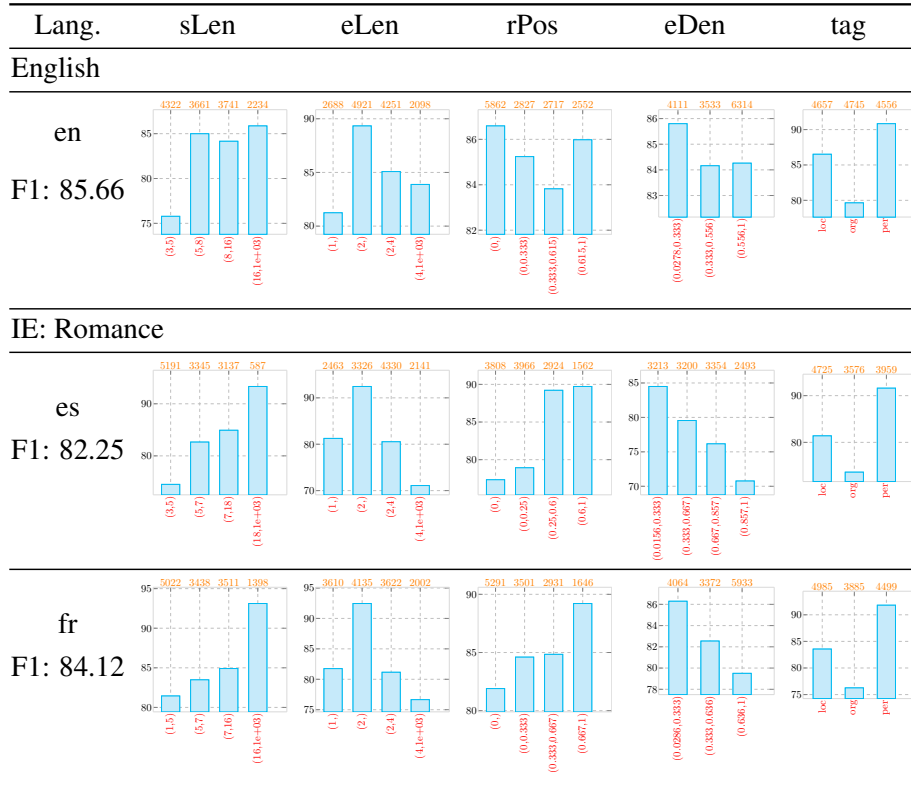


Table 27: The single system diagnosis of ERNIE-M on WikiANN-NER task in *en*, *es* and *fr*. The first row shows the fine-grained performance on English test data, and the second and third rows show the zero-shot transfer fine-grained performance on languages in the IE: Romance language family.

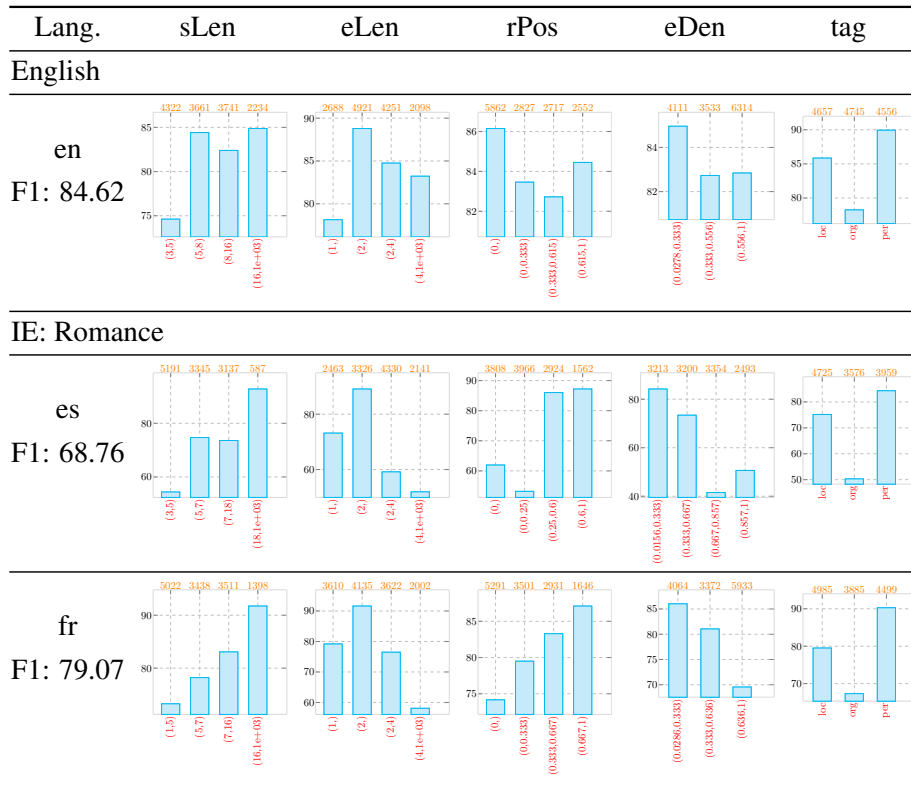


Table 28: The single system diagnosis of XLM-R on WikiANN-NER task in *en*, *es* and *fr*. The first row shows the fine-grained performance on English test data, and the second and third rows show the zero-shot transfer fine-grained performance on languages in the IE: Romance language family.

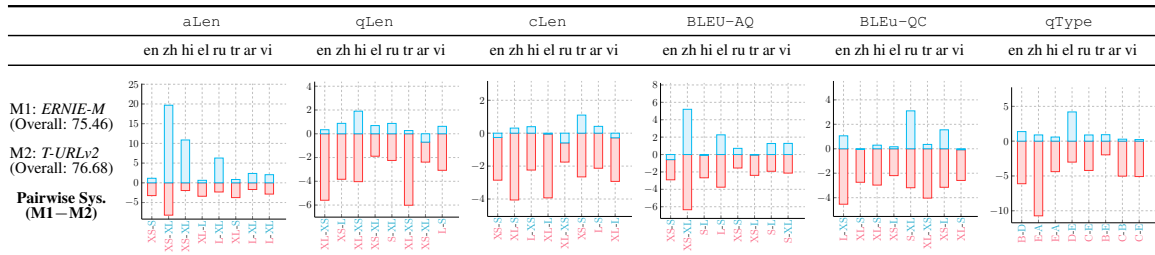


Table 29: Pairwise system diagnosis of ERNIE-M and T-URLv2 for XQuAD task. “M1–M2” represents the performance difference between M1 and M2. We classify the attribute values into four *categories*: extra-small (XS), small (S), large (L) and extra-large (XL) values. The top 5 most frequent question types (q_{Type}) include what (A), how (B), who (C), when (D) and which (E), ranked by their frequencies in training set. In the *pairwise system diagnosis* histogram, blue (red) x ticklabels represents the bucket value of a specific attribute on which system M1 surpasses (under-performs) M2 by the largest margin that is illustrated by a blue (red) bin. The blue-only x ticklabels (e.g., -D) indicate that M1 outperforms M2 in all categories of an attribute.