

What is “Typological Diversity” in NLP?

Esther Ploeger*[✳] Wessel Poelman*[✳] Miryam de Lhoneux*[✳] Johannes Bjerva*[✳]

[✳]Department of Computer Science, Aalborg University, Denmark

[✳]Department of Computer Science, KU Leuven, Belgium

{espl, jbjerva}@cs.aau.dk {wessel.poelman, miryam.delhoneux}@kuleuven.be

Abstract

The NLP research community has devoted increased attention to languages beyond English, resulting in considerable improvements for multilingual NLP. However, these improvements only apply to a small subset of the world’s languages. An increasing number of papers aspires to enhance *generalizable* multilingual performance *across languages*. To this end, linguistic typology is commonly used to motivate language selection, on the basis that a broad typological sample ought to imply generalization across a broad range of languages. These selections are often described as being ‘typologically diverse’. In this meta-analysis, we systematically investigate NLP research that includes claims regarding typological diversity. We find there are no set definitions or criteria for such claims. We introduce metrics to approximate the diversity of resulting language samples along several axes and find that the results vary considerably across papers. Crucially, we show that skewed language selection can lead to overestimated multilingual performance. We recommend future work to include an operationalization of typological diversity that empirically justifies the diversity of language samples. To help facilitate this, we release the code for our diversity measures.¹

1 Introduction

Most research in the field of natural language processing (NLP) is conducted on the English language (Ruder et al., 2022). Competitive monolingual language modelling beyond English remains challenging, as current state-of-the-art methods rely on the availability of large amounts of data, which are not available for most other languages (Joshi et al., 2020). This data sparsity can be mitigated

* Equal contribution.

¹Our code and data are publicly available: <https://github.com/WPoelman/typ-div>




Figure 1 – There is an increase in the number of publications with ‘typological diversity’ claims across time.

by leveraging cross-lingual transfer through the training of a language model on multilingual data. Despite the potential of multilingual language modelling, methodologies are primarily developed for English. However, there is no guarantee that an approach that works well for one language will work equally well for others (Gerz et al., 2018). For instance, morphologically complex (or rich) languages can be over-segmented by current widely-used tokenization methods (Rust et al., 2021). Evaluation on a broad range of languages is important for drawing more generalizable conclusions about the performance of multilingual language technology (Bender, 2009, 2011; Pikuliak and Simko, 2022). Evaluating a tokenization method on only morphologically simple languages, such as English, can give an unrealistic image of the effectiveness of a tokenization method, because morphologically simple languages are generally easier to tokenize compared to complex ones.

Current work increasingly evaluates models on multiple languages, but because of resource constraints, it is not realistic to test a model on the thousands of languages in the world. In order to still ensure a degree of generalizability, previous work

recognizes the importance of diverse language sampling. [Ponti et al. \(2020\)](#) suggest that merely evaluating on a small set of similar languages is an unreliable method for estimating a multilingual model’s performance, since such evaluations lack robustness to previously unseen typological properties, for instance. In an attempt to address these sampling issues, NLP research increasingly claims to rely on linguistic typology (see [Figure 1](#)). Evaluating on a set of so-called ‘typologically diverse’ languages implies robustness over typological variations, implicitly claiming broad generalizability across languages. Importantly, however, there is a fundamental lack of common understanding of exactly what a claim of typological diversity entails.

In this work, we systematically investigate how the term ‘typological diversity’ is used in NLP research by surveying the literature for past and current practices. We find that there is no set definition of ‘typological diversity’ in NLP: justifications for claims range from (i) *none whatsoever*, to (ii) using language phylogeny as a proxy for typology, to (iii) basing diversity on a more or less well-founded subset of available typological features. As a consequence, typological language similarity and typological feature value inclusion vary considerably across papers (§5). Furthermore, we demonstrate why this is problematic, by showing how a skewed selection of evaluation languages can lead to overstated multilingual performance and poor generalizability (§6). We recommend future work using the term ‘typologically diverse’ to document their language samples, and to add a measure of typological diversity, such as the mean pairwise language distance (§5.1) or typological feature inclusion (§5.2). Better-documented language sampling is crucial for the development of multilingual NLP, because it provides insights into commonly overlooked linguistic difficulties of multilingual language modelling. We provide software to verify and visualize this, linked below the abstract.

2 Background

Language sampling techniques are used in linguistic typology as well as in NLP. Although these fields are both concerned with languages and generalizability, they typically have different objectives and resources, which we contrast in this section.

Objectives Language sampling has a long-standing history in linguistic typology. [Rijkhoff et al. \(1993\)](#) state that testing a linguistic hypoth-

esis on a few similar languages is not sufficient for drawing conclusions on human language as a whole. Rather, they argue that generalizable samples should capture the diversity of the world’s languages. Language samples can be constructed in a number of ways. [Rijkhoff and Bakker \(1998\)](#) distinguish three types of sampling: *random*, *probability* and *variety*. Random samples do not have any specific stratification criteria, whereas probability samples are constructed to minimize bias, for instance by avoiding a skew towards languages from particular language families or geographical areas. Variety samples are designed to contain as much linguistic variety as possible, including relatively uncommon typological characteristics ([Miestamo et al., 2016](#)). Finally, in addition to the three methods outlined, sampling can be based on data availability, which is defined as *convenience* sampling ([Velupillai, 2012](#), p. 50).

Language sampling methods in NLP have a similar goal: to sample languages in such a way that the findings (e.g., multilingual performance estimates) are expected to generalize across languages. Yet, their justifications are commonly less rigorous than those in linguistic typology. For language sampling methods in NLP, it is usually not explicitly mentioned which type of sampling is conducted, and for which reasons. In practice, most papers do not provide a justification for their selections (see §4.3). This is likely caused by data availability, albeit implicitly, which resembles convenience sampling. A rare exception to this is [Ponti et al. \(2020\)](#), who use variety sampling and mention that this ensures evaluation that is robust towards typological characteristics that were likely not seen during training.

Resources While the goal of language sampling for ensuring generalizable conclusions is similar between NLP and linguistic typology, their resources are typically different. In linguistic typology, the variables under investigation are often typological characteristics. A typologist may aim to test whether a correlation of features holds across languages. Such a correlation could then be referred to as a ‘universal’: languages with object-verb ordering would be postpositional ([Greenberg, 1963](#)), for instance. To avoid circularity, language sampling is not done directly based on those features. Instead, methods commonly rely on (combinations of) language family, genus, and geography ([Miestamo et al., 2016](#)).

By contrast, what is evaluated in multilingual NLP does not concern typological features directly, rather, model performance across languages. In fact, to support claims of generalizability, multilingual benchmarks must explicitly contain a range of typological phenomena (Bender, 2009, 2011; Piku-liak and Simko, 2022). Relying on geographical and genetic relations between languages is not guaranteed to ensure this. Stoll and Bickel (2013) argue that “*the kind of variables that define genealogical groups and tree shapes have a very different nature from the kind of variables that define typological diversity*”, based on work from Nichols (1996). For estimating the generalizability of a language sample in NLP evaluation, it is possible, and even desired, to have direct insight into the actual typological properties that are included in the selection (see also §6).

This motivation is similarly voiced by Samardzic et al. (2024), who introduce a data-driven method of relating language samples in particular datasets to a reference language sample that is assumed to be typologically diverse. Yet, our work differs in crucial ways. Our aims are to look at what it means to make *claims* of typological diversity in current research practices and how to characterize the language samples such claims are made about. Our methods work as standalone measures; we do not need a reference sample to compare against. In addition, our approach is centered around languages, thus only requiring language descriptions, not entire datasets. Samardzic et al. (2024) approximate ‘aggregated morphological features’ through the average word length in a corpus, but this number may vary across domains within the same language.

3 Systematic Review

We systematically investigate the use of the term ‘typological diversity’ in NLP research, by conducting a survey across relevant published work.

3.1 Data Collection

We retrieve papers from the entire ACL Anthology (*ACL, COLING, LREC, etc.) and from several top-tier AI venues where NLP research is published (NeurIPS, ICLR, AACL, etc.).² We select papers that contain the following search string in either their title or abstract:

```
typolog.+?div.+?|div.+?typolog.+?
```

²See Appendix A for details.

Our search string aims to capture a broad range of formulations, including: “typologically and genetically diverse languages” (Xu et al., 2020), “languages of diverse typologies” (Eskander et al., 2022) and “diverse languages in terms of language family and morphological typology” (Eskander et al., 2020b). Another commonly used phrase is “we evaluate on N typologically *different* languages”. While we acknowledge that some researchers might (unintentionally) use the two claims interchangeably, we only focus on claims regarding ‘diversity’. It is reasonable to state that languages are ‘typologically different’, even though they might be similar (e.g., Dutch and German).

3.2 Annotation Guidelines

The first two authors annotated the retrieved papers according to the following criteria:

1. Does the paper claim that a particular language set is typologically diverse?

If so, we annotate the following:

2. Does the paper introduce a dataset?
3. Which languages are included?

For criterion 1, we only consider claims made by the paper at hand. That is, if a paper uses a dataset which is claimed to be ‘typologically diverse’, but the current paper does not mention this, we do not consider this to be a claim. This also applies to related work: if a paper references a claim made by another paper, but they themselves do not make a claim, then we do not consider this as a claim. The reason that we annotate whether a paper introduces a new dataset is that datasets are often starting points for subsequent projects. If datasets are claimed to be ‘typologically diverse’, this claim is spread not only through the paper introducing it, but also the dataset itself.³

Lastly, we annotate which languages are included in the papers making a claim. These are normalized to ISO-639-3 codes if available. Where applicable, we only select the languages about which the claim is made. For example, Kann et al. (2020) and Shi et al. (2023) specifically distinguish between baseline languages and their typologically diverse test language selection.

³An overview of all papers introducing a dataset for ‘typologically diverse’ languages is in Appendix D, Table 3.

3.3 Inter-Annotator Agreement

In order to determine the agreement between the annotators, and estimate the quality of the annotation guidelines, we calculate Cohen’s κ , as there are two annotators and we treat the values as binary. The agreement values of all annotation items are shown in Table 1.

Annotation item	κ	Agreement
1. has_claim	0.78	Substantial
2. introduces_dataset	0.71	Substantial
3. iso_codes	0.44	Moderate

Table 1 – Inter-annotator agreement per item.

For the annotations concerning claims (1) and datasets (2), we achieve substantial agreement. We resolve the disagreements by discussing the annotations and merging them into a single label. For the sake of transparency, and to support incentives for human label variation (Plank, 2022), we also release the non-aggregated annotations.⁴ The agreement for the language selection annotation (3) is somewhat low, which is partially explained by the fact that we calculate agreement over all languages together, rather than on the micro-level for individual languages. This is because there is no ground truth regarding the number of languages to annotate. Furthermore, a number of inconsistencies are due to ISO-639-3 variants (jap → jpn or ger → deu) and different codes for possibly ambiguous language mentions (‘Norwegian’ could be nor or nob). We iteratively resolve these issues to get our final list of languages per paper.




Figure 2 – Number of papers with a claim by venue.

⁴At the repository listed below the abstract.

4 Results

4.1 Publication Overview

In total we retrieve 194 papers, of which 110 are found to contain a claim. The rest of this analysis is based on these 110 papers. Figure 2 shows the most common venues where papers with a claim have been published. The top six venues are all high-ranking NLP conferences, with SIGMORPHON being the first workshop at the seventh spot. The AI-centered venues only have a small number of papers with claims; ICLR and AAIi both have two, ICML one, and the rest none. In total, 38 papers introduce a dataset, most of which are published at ACL and LREC since 2020 (Figure 1 and 2).⁵

4.2 Language Overview




Figure 3 – Papers by number of used languages.

Four of the papers that claim typological diversity do not mention which languages they use at all. The other papers use a median of 11 languages ($Q1 = 8$, $Q3 = 18$), with a minimum of 2 and a maximum of 90. In total, the papers use 315 unique languages, of which 160 are used just once. Figure 3 shows the distribution of the number of languages used per paper.

English is the most used language (63 papers), followed by German (60), Russian (58) and Finnish (57).⁶ The paper that uses the most languages is Vylomova et al. (2020), which contains 90. Figure 4 shows where these languages are primarily

⁵Details are listed in Appendix D, Table 4.

⁶The long tail of this distribution is partially included in Appendix C, Figure 8.

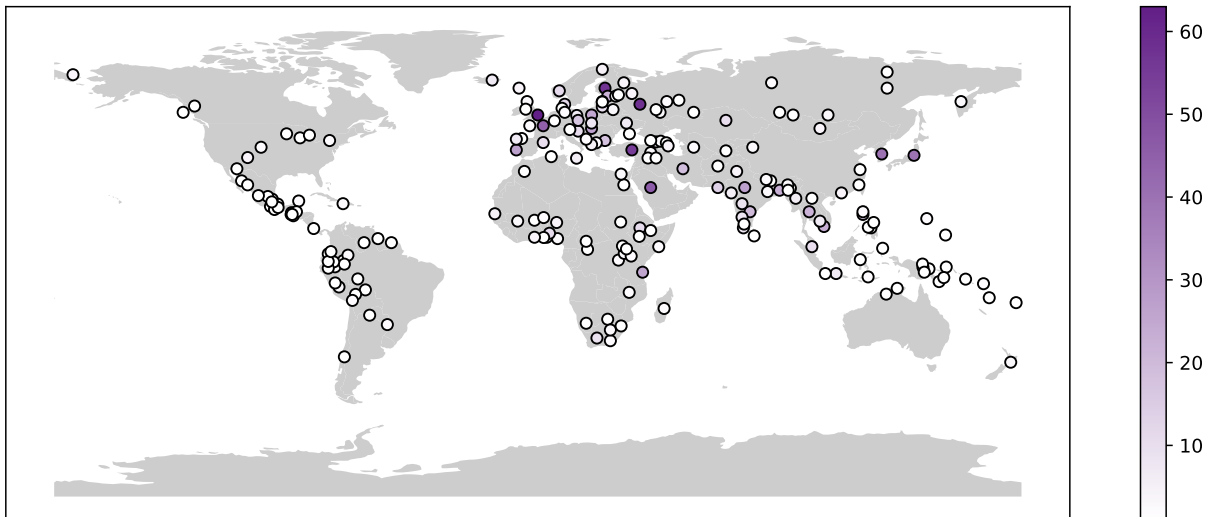


Figure 4 – Map of languages in all papers claiming ‘typological diversity’, where the hue corresponds number of papers that uses a language. Coordinates are taken from WALS.

spoken according to the World Atlas of Language Structures (WALS) (Dryer and Haspelmath, 2013). We observe a skew towards languages spoken in Europe, where multiple languages are used in more than 50 papers, while not a single language from the Americas is included in more than 10 papers.

4.3 Justifications

The papers include a wide spectrum of justifications for their claims. A large portion provide no justification at all. Some use genealogy: the selection of 24 languages from Xu et al. (2022) aims to cover “a reasonable variety of language families”, while the dataset created by Zhang et al. (2023) consists of “[18] languages (...) from 10 language families and 13 sub-families”. Others use a selection of typological features, for instance, Mott et al. (2020) mention that “the nine languages in our corpus cover five primary language families (...), and a range of morphological phenomena (...)”. Some also mention typological databases in their language selection: Muradoglu and Hulden (2022) consider “languages that exhibit varying degrees of complexity for inflection. We also consider morphological characteristics coded in WALS (...)”. A rather systematic approach to language selection is found in Jancso et al. (2020). They use a clustering algorithm on vectors with features from two typological databases to find the most distant clusters to sample languages from.

Still, most other papers justify their typological diversity in a ‘post-hoc’ way. They generally do not mention *initial* language selection considerations.

Rather, they mention how diverse the sample is *after* it has been created. The XTREME-R dataset (Ruder et al., 2021), which we will discuss further in §6, is exemplary of this. They mention ‘diversity indices’ that cover family, as well as typology and compare theirs to other datasets, but do not detail how or if these were used in sampling.

5 Typological Analysis

From Section 4.2, it follows that there is a geographical skew, where languages from certain areas are over-represented (Figure 4). This is not surprising, as it is often more feasible to select languages for which there are existing resources and typological descriptions, than languages for which these have to be gathered from scratch. This constitutes bibliographic bias (Rijkhoff and Bakker, 1998). However, we cannot draw conclusions about typological diversity on this basis, as typology is not equal to geography, nor phylogeny (Cysouw, 2013). Therefore, in this section, we analyze the typological diversity according to metrics that take into account syntactic features, and the absolute included typological feature values. Ponti et al. (2020) previously quantified language diversity with *typology* (URIEL), *geography*, and *family* indices. These metrics do not take into account the distance between pairs of languages, and thus do not provide insights into the independence of languages in a sample. Additionally, no code or clear descriptions are provided, which limits their use.




Figure 5 – Distributions of mean pairwise lang2vec distances and feature inclusion per paper. On the **left** are approximations based on common justifications for claiming ‘typological diversity’: geography ($\mu = 0.28$, $\sigma = 0.11$) and genealogy ($\mu = 0.94$, $\sigma = 0.05$). On the **right** two different approximations based on typological features: MPSD ($\mu = 0.64$, $\sigma = 0.07$) and Grambank feature value inclusion ($\mu = 0.72$, $\sigma = 0.17$).

Our metrics⁷ include pairwise language distances (§5.1) and absolute typological feature value inclusion (§5.2) across papers. Lastly, we look into the relationship between the number of included languages and typological diversity, to help guide future language selection efforts (§5.3).

5.1 Mean Language Distance

We first approximate the typological diversity of each paper’s language selection with lang2vec (Littell et al., 2017). This toolkit provides language distances that are calculated based on language vectors from URIEL, a resource that contains information from a range of typological databases, including WALS. We choose lang2vec specifically, because its aggregated information enables analysis across databases, and because its vectorized format mitigates issues stemming from incomplete feature coverage in typological databases.

For each paper in our survey that claims a typologically diverse language set and that is covered in lang2vec (312 / 315 unique languages), we calculate the MEAN PAIRWISE SYNTACTIC DISTANCE (MPSD) for the included languages. That is, we take the average of all pre-computed syntactic lang2vec distances for each pair of languages in a given paper’s selection. We only take into account languages that have at least 5% coverage in the URIEL vectors. We do this to make sure the very low coverage languages do not distort the results. This affects 31 papers in our dataset, resulting in a total of 242 unique languages. The top-right plot in Figure 5 shows the distribution of MPSDs across papers.⁸ We compare our typology-based approximations with mean geographical and

genetic pairwise distances. Interestingly, the genetic distance is much higher and less spread out than the syntactic distance. This emphasizes that genealogical sampling does not by default ensure typological diversity.

The typological distances vary considerably, with outliers on either side (Figure 5). In papers claiming typological diversity, the minimum MPSD is found in Goel et al. (2022), who use English, French, and Spanish. The maximum MPSD is found for Vania et al. (2019) their selection of North Sámi, Galician, and Kazah.

5.2 Inclusion of Typological Features

Because lang2vec distance calculations are based on feature vectors, our previous analysis does not provide information regarding *which* typological feature values are actually included in the language set. For example, it does not tell us which word order variations are covered in a given language selection. This information can be useful though, because covering more typological feature values in an evaluation set means that robustness towards previously unseen typological characteristics is increased. Following the *saturation* measure by Mistamo et al. (2016), we look into the inclusion of individual typological feature values per language selection from each paper. We use the Grambank⁹ (Skirgård et al., 2023b) database, because it has high coverage for grammatical features and is currently actively maintained. For each feature in Grambank, we count whether all possible, non-missing feature values are represented in the paper’s language selection.¹⁰ That is, for a feature such as *GB020: Are there definite or specific articles?*, we count whether all non-missing options are represented by the particular set of languages

⁷Formal definitions are provided in Appendix B.

⁸To illustrate this: the distance between Danish and Norwegian (same family, same genus) is 0.22, Danish-Spanish (same family, different genus) is 0.59 and Danish-Japanese (different family, different genus) is 0.69.

⁹<https://grambank.clld.org/>

¹⁰Note that for our purposes, we treat ? as a missing value.

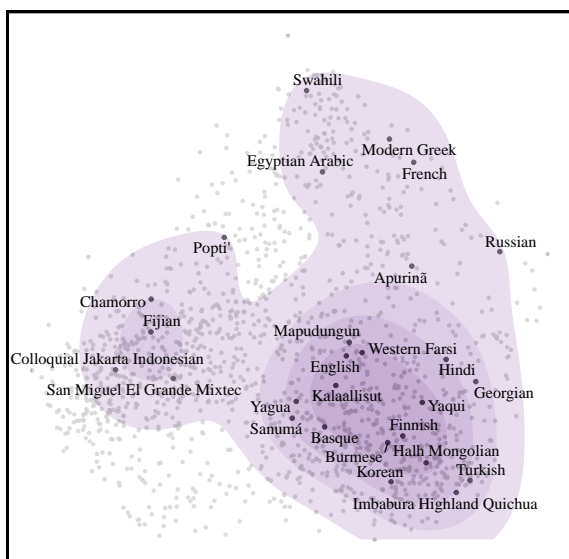


Figure 6 – PCA plot for the paper with the highest absolute feature coverage (Gutierrez-Vasques et al., 2021), where each point represents a language in Grambank.

from a paper and divide this by the total number of features. The spread of these values is shown in the bottom right graph in Figure 5. We treat missing languages the same as in the MPSD calculations; we discard them during calculation, but we do not discard the entire paper. Twelve papers are affected by this. The average Grambank feature value inclusion is 0.73.

While the accumulation of absolute feature values covered in a language selection provides useful robustness insights, it does not capture that languages consist of *combinations* of features that are not entirely independent. Skirgård et al. (2023a) show how language families can reliably be visualized as distinct groupings in Grambank’s typological design space using principle component analysis (PCA). Similarly, in Figure 6 we show the PCA plot for the paper that has the highest Grambank feature inclusion of all papers: Gutierrez-Vasques et al. (2021). While their language sample covers much of the Grambank’s typological space, we observe a skew towards certain languages (bottom right), and a lack of representation for others (e.g., bottom left). This means that even the paper with the highest feature value inclusion in our survey contains challenges when it comes to complete and fair evaluation.

5.3 Number of Languages

As shown in Figure 3, the number of included languages varies across papers. Should one aim to




Figure 7 – MPSD (a) and Grambank feature value inclusion (b) per paper by number of languages.

evaluate on as many languages as possible? What is the effect of selecting many (similar) languages? Figure 7a and 7b show the relationship between the number of languages and the MPSDs, and the absolute typological feature value inclusion, respectively. We find that smaller language sets (0-10 languages) exhibit considerable variation in terms of average syntactic distances, while the larger language sets (30+ languages) seem consistent between 0.6 and 0.7. This implies that adding more languages does not necessarily raise the average syntactic distance much. In NLP, adding more languages typically means adding more *similar* languages, since it is easier to incorporate existing datasets than it is to create new ones. As a result, the average performance will also be skewed towards these similar languages in evaluation (see §6). Furthermore, we observe that up to a certain point, including more languages implies that more typological feature values are covered. However, this increase flattens at approximately 40 languages. Importantly, this highlights the fact that simply adding more languages to an experimental study is not by itself a contribution in terms of typological generalizability – one must take care *which* languages are included. To illustrate this, we highlight the MKQA dataset (Longpre et al., 2021). This dataset includes 25 languages, more than twice the median, but has a rather average MPSD of 0.61 (and a Grambank feature inclusion of 0.89). This is in part due to its inclusion of several typologically similar languages: {English, German, Dutch}, {Cantonese, Mandarin}, {Spanish, Portuguese} and {Swedish, Danish, Norwegian}.

6 Downstream Implications

Inconsistencies in ‘typologically diverse’ language selection can have an effect on downstream evaluation. While reporting all metrics for all lan-

Subtask	Model	Overall	By F	Δ	Strong Pre	Weak Pre	Equal Pre & Suf	Strong Suf	Weak Suf	Little Aff	NA
Mewsl-X [★]	XML-R-L	45.75 (11)	36.23 (11)	-9.52	- (0)	- (0)	- (0)	47.86 (10)	24.60 (1)	- (0)	- (0)
	mBERT	38.58 (11)	27.29 (11)	-11.29	- (0)	- (0)	- (0)	41.09 (10)	13.50 (1)	- (0)	- (0)
XNLI [★]	XML-R	79.24 (15)	76.54 (15)	-2.70	- (0)	71.20 (1)	- (0)	80.06 (12)	- (0)	78.35 (2)	- (0)
	mBERT	66.51 (15)	60.17 (15)	-6.35	- (0)	49.30 (1)	- (0)	68.60 (12)	- (0)	62.60 (2)	- (0)
	mT5	84.85 (15)	82.92 (15)	-1.92	- (0)	80.60 (1)	- (0)	85.57 (12)	- (0)	82.60 (2)	- (0)
LAReQA [★]	XML-R-L	40.75 (11)	40.54 (11)	-0.22	- (0)	- (0)	- (0)	40.88 (9)	- (0)	40.20 (2)	- (0)
	mBERT	21.58 (11)	19.24 (11)	-2.35	- (0)	- (0)	- (0)	22.92 (9)	- (0)	15.55 (2)	- (0)
XQuAD [★]	XML-R-L	77.21 (11)	77.24 (11)	+0.04	- (0)	- (0)	- (0)	77.19 (9)	- (0)	77.30 (2)	- (0)
	mBERT	65.05 (11)	61.84 (11)	-3.21	- (0)	- (0)	- (0)	66.89 (9)	- (0)	56.80 (2)	- (0)
	mT5	81.54 (11)	80.55 (11)	-0.99	- (0)	- (0)	- (0)	82.10 (9)	- (0)	79.00 (2)	- (0)
MLQA [★]	XML-R-L	72.71 (7)	73.33 (7)	+0.62	- (0)	- (0)	- (0)	72.47 (6)	- (0)	74.20 (1)	- (0)
	mBERT	61.30 (7)	60.84 (7)	-0.46	- (0)	- (0)	- (0)	61.48 (6)	- (0)	60.20 (1)	- (0)
	mT5	75.59 (7)	75.97 (7)	+0.38	- (0)	- (0)	- (0)	75.43 (6)	- (0)	76.50 (1)	- (0)
Tatoeba [★]	XML-R	77.29 (41)	64.92 (36)	-12.36	- (0)	31.30 (1)	58.60 (1)	82.10 (28)	76.37 (3)	77.43 (3)	63.74 (5)
	mBERT	43.33 (41)	32.03 (36)	-11.30	- (0)	12.10 (1)	31.00 (1)	49.24 (28)	39.27 (3)	32.90 (3)	27.68 (5)
UD-POS [★]	XML-R-L	74.96 (38)	71.12 (36)	-3.84	- (0)	- (0)	74.30 (1)	79.75 (28)	71.05 (2)	45.98 (5)	84.50 (2)
	mBERT	70.90 (38)	64.43 (36)	-6.47	- (0)	- (0)	59.30 (1)	75.51 (28)	60.75 (2)	48.66 (5)	77.95 (2)
XCOPA [★]	XML-R	69.22 (11)	65.93 (9)	-3.28	- (0)	61.80 (1)	- (0)	73.93 (6)	- (0)	75.30 (2)	52.70 (2)
	mBERT	56.05 (11)	54.75 (9)	-1.30	- (0)	52.20 (1)	- (0)	57.70 (6)	- (0)	56.20 (2)	52.90 (2)
	mT5	74.89 (11)	73.24 (9)	-1.65	- (0)	74.10 (1)	- (0)	78.00 (6)	- (0)	77.60 (2)	63.25 (2)
WikiANN-NER [★]	XML-R-L	64.43 (48)	62.02 (40)	-2.41	- (0)	69.90 (1)	62.10 (1)	66.92 (31)	61.37 (3)	48.17 (4)	63.66 (8)
	mBERT	62.68 (48)	61.73 (40)	-0.95	- (0)	72.70 (1)	65.00 (1)	64.93 (31)	57.23 (3)	49.38 (4)	61.12 (8)
TyDiQA [★]	XML-R-L	64.29 (9)	62.57 (8)	-1.72	- (0)	66.40 (1)	- (0)	65.67 (6)	- (0)	59.10 (1)	59.10 (1)
	mBERT	58.36 (9)	55.09 (8)	-3.26	- (0)	59.70 (1)	- (0)	60.97 (6)	- (0)	46.20 (1)	53.50 (1)
	mT5	81.94 (9)	83.73 (8)	+1.78	- (0)	87.20 (1)	- (0)	80.52 (6)	- (0)	83.60 (1)	83.60 (1)

Table 2 – XTREME-R results grouped by inflection type (WALS 26A). **Overall** refers to the average over all languages. **By Feature** is the average of the WALS feature averages, excluding languages for which there is no coverage. The delta shows the difference of **Overall** and **By F**. Tasks that have coverage for all their included languages are on top, those partially covered are in the bottom portion. Morphological inflection types are: *Aff* = affixing, *Suf* = suffixing, *Pre* = prefixing and *NA* referring to *Not Available* (in WALS). The **highest** and the **lowest** non-zero number of languages per grouping are highlighted. The (*italicized*) number refers to the number of languages in a particular subset. Metrics: \blacklozenge = Accuracy, \star = mAP@20, \spadesuit = F1.

guages separately is preferable, space constraints often lead to reporting *averages* across languages in multilingual evaluation. However, Anastasopoulos (2019) demonstrates that when a language set contains a skew towards a certain language family, simply taking the micro average over all languages for evaluation gives an overestimation of performance. While family can serve as a proxy for typological diversity, we here present an analysis based directly on typological characteristics. In this way, we gain more fine-grained insights into the effects of typologically skewed language selection for model evaluation.

Specifically, we use XTREME-R (Ruder et al., 2021), a popular¹¹ ‘typologically diverse’ multilingual benchmark covering several tasks and domains. XTREME-R contains and expands upon ten existing datasets. As a result, not all tasks are available in all languages; while there are 50 unique languages in total, the number of languages per task

ranges from 7 to 48.¹² We group these languages by the ‘Prefixing vs. Suffixing in Inflectional Morphology’ (26A) and ‘Order of Subject, Object and Verb’ (81A) features, as provided by WALS (Dryer, 2013b,a). These features have high language coverage, which makes our analysis as comprehensive as possible. Table 2 shows groupings and coverage of the inflectional feature (26A). The word order analysis is included in Appendix E, Table 5.

Firstly, we see that XTREME-R contains a skew in terms of morphological inflection. For all tasks, the majority of languages are *strongly suffixing* (see orange cells). Some feature values, such as *equal prefixing and suffixing* and *weak prefixing* are underrepresented, with at most one language per subtask. Remarkably, *strong prefixing* does not appear in XTREME-R at all. This implies that one should be careful with the *implication* of generalizability that evaluating on a ‘typologically diverse’ dataset gives.

The delta column shows the difference between the macro average over all languages and the macro average per feature value. Here, we observe simi-

¹¹The original XTREME dataset has almost 800 citations as of writing, XTREME-R, which is an extension of the original, more than 100.

¹²See Appendix F, Figure 9.

lar patterns as in [Anastasopoulos \(2019\)](#), namely that estimations of multilingual model performance vary considerably when accounting for typological imbalances. We suggest that drawing and implying generalizable conclusions about typologically diverse evaluation should ideally be supported by showing that performance holds across a range of typological properties.

7 Insights and Recommendations

We systematically analyze claims of ‘typological diversity’ in NLP research. We approximate diversity in terms of average syntactic language distance (MPSD) and absolute typological feature value inclusion in Grambank. Our analysis shows that (1) there is no consistent definition or methodology when making ‘typological diversity’ claims, (2) estimates of typological diversity exhibit a considerable variation across papers and (3) aggregated results can give distorted views of multilingual performance estimates.

We recommend future approaches to include an operationalization of ‘typological diversity’ when making such claims. This can be in terms of generalizability or regarding a particular typological phenomenon of interest. Additionally, we recommend adding an empirical justification, especially when claims relate to generalizability. The feature value inclusion, MPSD or PCA plots we have shown are examples of these justifications. Including these has the potential to benefit multilingual NLP, as it enables more fine-grained insights into typological challenges of language modelling.

Limitations

Our measures of typological diversity are approximations. There is no typological database that covers every aspect of every language. This also ties into the bibliographic bias of typological databases in general. In addition, languages are not necessarily as discrete as they appear in typological resources either ([Levshina et al., 2023](#); [Baylor et al., 2024](#)). Nonetheless, we believe that the reporting of typological diversity can be more principled than it currently is, despite incomplete resources.

Another limitation is that our search is based only on the title and abstract of papers. There are almost certainly papers that contain a claim in other sections. However, we chose this route since we wanted to (1) cast a wide net, (2) perform some pre-filtering of papers, assuming that a mention of

‘typological diversity’ in the title or abstract is a good indicator that it is a prominent aspect of that paper, and (3) make sure that the data we could get was of good quality, which is a lot harder with (old) pdf files for example.

Our arguments for assessing linguistic typology directly in NLP, rather than phylogeny and geography as proxies, relates to claims of ‘typological diversity’. We acknowledge that geographical and genealogical information could be useful for other purposes.

Ethics Statement

In this paper, we investigate the use of the term ‘typological diversity’ in multilingual NLP research. We do not make a claim that NLP applications *should* be expanded to as many languages as possible, as introducing these technologies are not necessarily always a positive influence ([Bird, 2020](#)). Instead, our aim is to gain a common understanding and clear explanations regarding the usage of these claims.

Acknowledgements

EP and JB are funded by the Carlsberg Foundation, under the Semper Ardens: Accelerate programme (project nr. CF21-0454). WP is funded by a KU Leuven Bijzonder Onderzoeksfonds C1 project with reference C14/23/096.

We thank the TypNLP group at Aalborg University and the LAGoM-NLP group at KU Leuven for insightful feedback on earlier versions of this paper, in particular Heather Lent. We also thank Thomas Bauwens for his help in creating the tables (using `fiject`).

References

- Antonis Anastasopoulos. 2019. [A note on evaluating multilingual benchmarks](#). Accessed: 2024-01-09.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. [On the cross-lingual transferability of monolingual representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4623–4637, 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 Association for Computational Linguistics*, 7:597–610.

- Akari Asai, Shayne Longpre, Jungo Kasai, Chia-Hsuan Lee, Rui Zhang, Junjie Hu, Ikuya Yamada, Jonathan H. Clark, and Eunsol Choi. 2022. [MIA 2022 shared task: Evaluating cross-lingual open-retrieval question answering for 16 diverse languages](#). In *Proceedings of the Workshop on Multilingual Information Access (MIA)*, pages 108–120, Seattle, USA. Association for Computational Linguistics.
- Martijn Bartelds, Nay San, Bradley McDonnell, Dan Jurafsky, and Martijn Wieling. 2023. [Making more of little data: Improving low-resource automatic speech recognition using data augmentation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 715–729, Toronto, Canada. Association for Computational Linguistics.
- Emi Baylor, Esther Ploeger, and Johannes Bjerva. 2024. [Multilingual gradient word-order typology from Universal Dependencies](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 42–49, St. Julian’s, Malta. Association for Computational Linguistics.
- Emily M. Bender. 2009. [Linguistically naïve != language independent: Why NLP needs linguistic typology](#). In *Proceedings of the EACL 2009 Workshop on the Interaction between Linguistics and Computational Linguistics: Virtuous, Vicious or Vacuous?*, pages 26–32, Athens, Greece. Association for Computational Linguistics.
- Emily M. Bender. 2011. [On achieving and evaluating language-independence in NLP](#). *Linguistic Issues in Language Technology*, 6.
- Steven Bird. 2020. [Decolonising speech and language technology](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3504–3519, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Jan A. Botha, Zifei Shan, and Daniel Gillick. 2020. [Entity Linking in 100 Languages](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7833–7845, Online. Association for Computational Linguistics.
- Sven Buechel, Susanna Rucker, and Udo Hahn. 2020. [Learning and evaluating emotion lexicons for 91 languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1202–1217, Online. Association for Computational Linguistics.
- Emanuele Bugliarello, Fangyu Liu, Jonas Pfeiffer, Siva Reddy, Desmond Elliott, Edoardo Maria Ponti, and Ivan Vulić. 2022. [IGLUE: A benchmark for transfer learning across modalities, tasks, and languages](#). In *International Conference on Machine Learning*, pages 2370–2392. PMLR.
- 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](#). *Transactions of the Association for Computational Linguistics*, 8:454–470.
- 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 the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Michael Cysouw. 2013. [Disentangling geography from genealogy](#). In *Space in Language and Linguistics*, pages 21–37. De Gruyter.
- Cheikh M. Bamba Dione, David Ifeoluwa Adelani, Peter Nabende, Jesujoba Alabi, Thapelo Sindane, Happy Buzaaba, Shamsuddeen Hassan Muhammad, Chris Chinenye Emezue, Perez Ogayo, Anuoluwapo Aremu, Catherine Gitau, Derguene Mbaye, Jonathan Mukiibi, Blessing Sibanda, Bonaventure F. P. Dossou, Andiswa Bukula, Rooweither Mabuya, Allahsera Auguste Tapo, Edwin Munkoh-Buabeng, Victoire Memdjokam Koagne, Fatoumata Ouoba Kabore, Amelia Taylor, Godson Kalipe, Tebogo Macucwa, Vukosi Marivate, Tajuddeen Gwadabe, Mboning Tchiazee Elvis, Ikechukwu Onyenwe, Gratien Atindogbe, Tolulope Adelani, Idris Akinade, Olanrewaju Samuel, Marien Nahimana, Théogène Musabeyezu, Emile Niyomutabazi, Ester Chimhenga, Kudzai Gotosa, Patrick Mizha, Apelete Agbolo, Seydou Traore, Chinedu Uchechukwu, Aliyu Yusuf, Muhammad Abdullahi, and Dietrich Klakow. 2023. [MasakhaPOS: Part-of-speech tagging for typologically diverse African languages](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10883–10900, Toronto, Canada. Association for Computational Linguistics.
- Matthew S. Dryer. 2013a. [Order of subject, object and verb \(v2020.3\)](#). In Matthew S. Dryer and Martin Haspelmath, editors, *The World Atlas of Language Structures Online*. Zenodo.
- Matthew S. Dryer. 2013b. [Prefixing vs. suffixing in inflectional morphology \(v2020.3\)](#). In Matthew S. Dryer and Martin Haspelmath, editors, *The World Atlas of Language Structures Online*. Zenodo.
- Matthew S. Dryer and Martin Haspelmath, editors. 2013. [WALS Online \(v2020.3\)](#). Max Planck Institute for Evolutionary Anthropology, Leipzig.
- Ramy Eskander, Francesca Callejas, Elizabeth Nichols, Judith Klavans, and Smaranda Muresan. 2020a. [MorphAGram, evaluation and framework for unsupervised morphological segmentation](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 7112–7122, Marseille, France. European Language Resources Association.

- Ramy Eskander, Cass Lowry, Sujay Khandagale, Judith Klavans, Maria Polinsky, and Smaranda Muresan. 2022. [Unsupervised stem-based cross-lingual part-of-speech tagging for morphologically rich low-resource languages](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4061–4072, Seattle, United States. Association for Computational Linguistics.
- Ramy Eskander, Smaranda Muresan, and Michael Collins. 2020b. [Unsupervised cross-lingual part-of-speech tagging for truly low-resource scenarios](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4820–4831, Online. Association for Computational Linguistics.
- Jack FitzGerald, Christopher Hench, Charith Peris, Scott Mackie, Kay Rottmann, Ana Sanchez, Aaron Nash, Liam Urbach, Vishesh Kakarala, Richa Singh, Swetha Ranganath, Laurie Crist, Misha Britan, Wouter Leeuwis, Gokhan Tur, and Prem Nataraajan. 2023. [MASSIVE: A 1M-example multilingual natural language understanding dataset with 51 typologically-diverse languages](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4277–4302, Toronto, Canada. Association for Computational Linguistics.
- Daniela Gerz, Ivan Vulić, Edoardo Maria Ponti, Roi Reichart, and Anna Korhonen. 2018. [On the relation between linguistic typology and \(limitations of\) multilingual language modeling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 316–327, Brussels, Belgium. Association for Computational Linguistics.
- Michael Ginn, Sarah Moeller, Alexis Palmer, Anna Stacey, Garrett Nicolai, Mans Hulden, and Miikka Silfverberg. 2023. [Findings of the SIGMORPHON 2023 shared task on interlinear glossing](#). In *Proceedings of the 20th SIGMORPHON workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 186–201, Toronto, Canada. Association for Computational Linguistics.
- Goran Glavaš, Mladen Karan, and Ivan Vulić. 2020. [XHate-999: Analyzing and detecting abusive language across domains and languages](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6350–6365, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Anmol Goel, Charu Sharma, and Ponnurangam Kumaraguru. 2022. [An unsupervised, geometric and syntax-aware quantification of polysemy](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10565–10574, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Omer Goldman, Khuyagbaatar Batsuren, Salam Khalifa, Aryaman Arora, Garrett Nicolai, Reut Tsarfaty, and Ekaterina Vylomova. 2023. [SIGMORPHON–UniMorph 2023 shared task 0: Typologically diverse morphological inflection](#). In *Proceedings of the 20th SIGMORPHON workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 117–125, Toronto, Canada. Association for Computational Linguistics.
- Joseph Harold Greenberg. 1963. Some universals of grammar with particular reference to the order of meaningful elements. In *Universals of Language*, pages 40–70. MIT press, Cambridge, MA.
- Shivanshu Gupta, Yoshitomo Matsubara, Ankit Chadha, and Alessandro Moschitti. 2023. [Cross-lingual knowledge distillation for answer sentence selection in low-resource languages](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 14078–14092, Toronto, Canada. Association for Computational Linguistics.
- Ximena Gutierrez-Vasques, Christian Bentz, Olga Sozinova, and Tanja Samardzic. 2021. [From characters to words: the turning point of BPE merges](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3454–3468, Online. Association for Computational Linguistics.
- Leonhard Hennig, Philippe Thomas, and Sebastian Möller. 2023. [MultiTACRED: A multilingual version of the TAC relation extraction dataset](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3785–3801, Toronto, Canada. Association for Computational Linguistics.
- Peter Juel Henriksen and Marcus Uneson. 2012. [SMALLWorlds – multilingual content-controlled monologues](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 3362–3368, Istanbul, Turkey. European Language Resources Association (ELRA).
- Chia-Chien Hung, Anne Lauscher, Ivan Vulić, Simone Ponzetto, and Goran Glavaš. 2022. [Multi2WOZ: A robust multilingual dataset and conversational pre-training for task-oriented dialog](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3687–3703, Seattle, United States. Association for Computational Linguistics.
- Anna Jancso, Steven Moran, and Sabine Stoll. 2020. [The ACQDIV corpus database and aggregation pipeline](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 156–165, Marseille, France. European Language Resources Association.
- 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 the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Katharina Kann, Ophélie Lacroix, and Anders Søgaard. 2020. [Weakly Supervised POS Taggers Perform Poorly on Truly Low-Resource Languages](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8066–8073.
- Jordan Kodner, Salam Khalifa, Khuyagbaatar Batsuren, Hossep Dolatian, Ryan Cotterell, Faruk Akkus, Antonios Anastasopoulos, Taras Andrushko, Aryaman Arora, Nona Atanalov, Gábor Bella, Elena Budianskaya, Yustinus Ghanggo Ate, Omer Goldman, David Guriel, Simon Guriel, Silvia Guriel-Agiashvili, Witold Kieraś, Andrew Krizhanovsky, Natalia Krizhanovsky, Igor Marchenko, Magdalena Markowska, Polina Mashkovtseva, Maria Nepomniashchaya, Daria Rodionova, Karina Scheifer, Alexandra Sorova, Anastasia Yemelina, Jeremiah Young, and Ekaterina Vylomova. 2022. [SIGMORPHON–UniMorph 2022 shared task 0: Generalization and typologically diverse morphological inflection](#). In *Proceedings of the 19th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 176–203, Seattle, Washington. Association for Computational Linguistics.
- Natalia Levshina, Savithry Namboodiripad, Marc Allasonnière-Tang, Mathew Kramer, Luigi Talamo, Annemarie Verkerk, Sasha Wilmoth, Gabriela Garrido Rodriguez, Timothy Michael Gupton, Evan Kidd, Zoey Liu, Chiara Naccarato, Rachel Nordlinger, Anastasia Panova, and Natalia Stoyanova. 2023. [Why we need a gradient approach to word order](#). *Linguistics*, 61(4):825–883.
- 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.
- 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, Valencia, Spain. Association for Computational Linguistics.
- Fangyu Liu, Emanuele Bugliarello, Edoardo Maria Ponti, Siva Reddy, Nigel Collier, and Desmond Elliott. 2021. [Visually grounded reasoning across languages and cultures](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10467–10485, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Shayne Longpre, Yi Lu, and Joachim Daiber. 2021. [MKQA: A linguistically diverse benchmark for multilingual open domain question answering](#). *Transactions of the Association for Computational Linguistics*, 9:1389–1406.
- Olga Majewska, Ivan Vulić, Diana McCarthy, and Anna Korhonen. 2020. [Manual clustering and spatial arrangement of verbs for multilingual evaluation and typology analysis](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4810–4824, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Matti Miestamo, Dik Bakker, and Antti Arppe. 2016. [Sampling for variety](#). *Linguistic Typology*, 20(2):233–296.
- Justin Mott, Ann Bies, Stephanie Strassel, Jordan Kodner, Caitlin Richter, Hongzhi Xu, and Mitchell Marcus. 2020. [Morphological segmentation for low resource languages](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 3996–4002, Marseille, France. European Language Resources Association.
- Saliha Muradoglu and Mans Hulden. 2022. [Eeny, meeny, miny, moe. how to choose data for morphological inflection](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7294–7303, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Johanna Nichols. 1996. The comparative method as heuristic. In Mark Durie, editor, *The Comparative Method Reviewed*, chapter 3, pages 39–71. Oxford University Press, New York.
- Joakim Nivre, Rogier Blokland, Niko Partanen, and Michael Rießler. 2018. Universal dependencies 2.2.
- Frank Palma Gomez, Subhadarshi Panda, Michael Flor, and Alla Rozovskaya. 2023. [Using neural machine translation for generating diverse challenging exercises for language learner](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6115–6129, Toronto, Canada. Association for Computational Linguistics.
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. [Cross-lingual name tagging and linking for 282 languages](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1946–1958, Vancouver, Canada. Association for Computational Linguistics.

- Jonas Pfeiffer, Gregor Geigle, Aishwarya Kamath, Jan-Martin Steitz, Stefan Roth, Ivan Vulić, and Iryna Gurevych. 2022. [xGQA: Cross-lingual visual question answering](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2497–2511, Dublin, Ireland. Association for Computational Linguistics.
- Matúš Pikuliak and Marian Simko. 2022. [Average is not enough: Caveats of multilingual evaluation](#). In *Proceedings of the The 2nd Workshop on Multilingual Representation Learning (MRL)*, pages 125–133, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Barbara Plank. 2022. [The “problem” of human label variation: On ground truth in data, modeling and evaluation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10671–10682, Abu Dhabi, United Arab Emirates. 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 common-sense reasoning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2362–2376, Online. Association for Computational Linguistics.
- Jan Rijkhoff and Dik Bakker. 1998. [Language sampling](#). *Linguistic Typology*, 2(3):263–314.
- Jan Rijkhoff, Dik Bakker, Kees Hengeveld, and Peter Kahrel. 1993. [A method of language sampling](#). *Studies in Language. International Journal sponsored by the Foundation “Foundations of Language”*, 17(1):169–203.
- Uma Roy, Noah Constant, Rami Al-Rfou, Aditya Barua, Aaron Phillips, and Yinfei Yang. 2020. [LAREQA: 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.
- Sebastian Ruder, Noah Constant, Jan Botha, Aditya Siddhant, Orhan Firat, Jinlan Fu, Pengfei Liu, Junjie Hu, Dan Garrette, Graham Neubig, and Melvin Johnson. 2021. [XTREME-R: Towards more challenging and nuanced multilingual evaluation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10215–10245, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Sebastian Ruder, Ivan Vulić, and Anders Søgaard. 2022. [Square one bias in NLP: Towards a multi-dimensional exploration of the research manifold](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2340–2354, Dublin, Ireland. Association for Computational Linguistics.
- Phillip Rust, Jonas Pfeiffer, Ivan Vulić, Sebastian Ruder, and Iryna Gurevych. 2021. [How good is your tokenizer? on the monolingual performance of multilingual language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3118–3135, Online. Association for Computational Linguistics.
- Tanja Samardzic, Ximena Gutierrez, Christian Bentz, Steven Moran, and Olga Pelloni. 2024. [A measure for transparent comparison of linguistic diversity in multilingual NLP data sets](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3367–3382, Mexico City, Mexico. Association for Computational Linguistics.
- Farhan Samir and Miikka Silfverberg. 2023. [Understanding compositional data augmentation in typologically diverse morphological inflection](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 277–291, Singapore. Association for Computational Linguistics.
- Gabriele Sarti, Arianna Bisazza, Ana Guerberof-Arenas, and Antonio Toral. 2022. [DivEMT: Neural machine translation post-editing effort across typologically diverse languages](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7795–7816, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Felix Sasaki, Claudia Wegener, Andreas Witt, Dieter Metzger, and Jens Pöninghaus. 2002. [Co-reference annotation and resources: A multilingual corpus of typologically diverse languages](#). In *Proceedings of the Third International Conference on Language Resources and Evaluation (LREC’02)*, Las Palmas, Canary Islands - Spain. European Language Resources Association (ELRA).
- Freda Shi, Mirac Suzgun, Markus Freitag, Xuezhi Wang, Suraj Srivats, Soroush Vosoughi, Hyung Won Chung, Yi Tay, Sebastian Ruder, Denny Zhou, Dipanjan Das, and Jason Wei. 2023. [Language models are multilingual chain-of-thought reasoners](#). In *The Eleventh International Conference on Learning Representations (ICLR)*.
- Hedvig Skirgård, Hannah J. Haynie, Damián E. Blasi, Harald Hammarström, Jeremy Collins, Jay J. Latache, Jakob Lesage, Tobias Weber, Alena Witzlack-Makarevich, Sam Passmore, Angela Chira, Luke Maurits, Russell Dinnage, Michael Dunn, Ger Reesink, Ruth Singer, Claire Bowers, Patience Epps, Jane Hill, Outi Vesakoski, Martine Robbeets, Noor Karolin Abbas, Daniel Auer, Nancy A. Bakker, Giulia Barbos, Robert D. Borges, Swintha Danielsen, Luise Dorenbusch, Ella Dorn, John Elliott, Gida Falcone, Jana Fischer, Yustinus Ghanggo Ate, Hannah Gibson, Hans-Philipp Göbel, Jemima A. Goodall, Victoria Gruner, Andrew Harvey, Rebekah Hayes, Leonard Heer, Roberto E. Herrera

- Miranda, Nataliia Hübler, Biu Huntington-Rainey, Jessica K. Ivani, Marilen Johns, Erika Just, Eri Kashima, Carolina Kipf, Janina V. Klingenberg, Nikita König, Aikaterina Koti, Richard G. A. Kowalik, Olga Krasnoukhova, Nora L. M. Lindvall, Mandy Lorenzen, Hannah Lutzenberger, Tânia R. A. Martins, Celia Mata German, Suzanne van der Meer, Jaime Montoya Samamé, Michael Müller, Saliha Muradođlu, Kelsey Neely, Johanna Nickel, Miina Norvik, Cheryl Akinyi Oluoch, Jesse Peacock, India O. C. Pearey, Naomi Peck, Stephanie Petit, Sören Pieper, Mariana Poblete, Daniel Prestipino, Linda Raabe, Amna Raja, Janis Reimringer, Sydney C. Rey, Julia Rizaew, Eloisa Ruppert, Kim K. Salmon, Jill Sammet, Rhiannon Schembri, Lars Schlabach, Frederick W. P. Schmidt, Amalia Skilton, Wikaliler Daniel Smith, Hilário de Sousa, Kristin Sverredal, Daniel Valle, Javier Vera, Judith Voß, Tim Witte, Henry Wu, Stephanie Yam, Jingting Ye, Maisie Yong, Tessa Yuditha, Roberto Zariquiey, Robert Forkel, Nicholas Evans, Stephen C. Levinson, Martin Haspelmath, Simon J. Greenhill, Quentin D. Atkinson, and Russell D. Gray. 2023a. [Grambank reveals the importance of genealogical constraints on linguistic diversity and highlights the impact of language loss](#). *Science Advances*, 9(16):eadg6175.
- Hedvig Skirgård, Hannah J. Haynie, Harald Hammarström, Damián E. Blasi, Jeremy Collins, Jay Latache, Jakob Lesage, Tobias Weber, Alena Witzlack-Makarevich, Michael Dunn, Ger Reesink, Ruth Singer, Claire Bowern, Patience Epps, Jane Hill, Outi Vesakoski, Noor Karolin Abbas, Sunny Ananth, Daniel Auer, Nancy A. Bakker, Giulia Barbos, Anina Bolls, Robert D. Borges, Mitchell Browen, Lennart Chevallier, Swintha Danielsen, Sinoël Dohlen, Luise Dorenbusch, Ella Dorn, Marie Duhamel, Farah El Haj Ali, John Elliott, Giada Falcone, Anna-Maria Fehn, Jana Fischer, Yustinus Ghanggo Ate, Hannah Gibson, Hans-Philipp Göbel, Jemima A. Goodall, Victoria Gruner, Andrew Harvey, Rebekah Hayes, Leonard Heer, Roberto E. Herrera Miranda, Nataliia Hübler, Biu H. Huntington-Rainey, Guglielmo Inglese, Jessica K. Ivani, Marilen Johns, Erika Just, Ivan Kapitonov, Eri Kashima, Carolina Kipf, Janina V. Klingenberg, Nikita König, Aikaterina Koti, Richard G. A. Kowalik, Olga Krasnoukhova, Kate Lynn Lindsey, Nora L. M. Lindvall, Mandy Lorenzen, Hannah Lutzenberger, Alexandra Marley, Tânia R. A. Martins, Celia Mata German, Suzanne van der Meer, Jaime Montoya, Michael Müller, Saliha Muradođlu, HunterGatherer, David Nash, Kelsey Neely, Johanna Nickel, Miina Norvik, Bruno Olson, Cheryl Akinyi Oluoch, David Osgarby, Jesse Peacock, India O.C. Pearey, Naomi Peck, Jana Peter, Stephanie Petit, Sören Pieper, Mariana Poblete, Daniel Prestipino, Linda Raabe, Amna Raja, Janis Reimringer, Sydney C. Rey, Julia Rizaew, Eloisa Ruppert, Kim K. Salmon, Jill Sammet, Rhiannon Schembri, Lars Schlabach, Frederick W. P. Schmidt, Dineke Schokkin, Jeff Siegel, Amalia Skilton, Hilário de Sousa, Kristin Sverredal, Daniel Valle, Javier Vera, Judith Voß, Daniel Wikalier Smith, Tim Witte, Henry Wu, Stephanie Yam, Jingting Ye, Maisie Yong, Tessa Yuditha, Roberto Zariquiey, Robert Forkel, Nicholas Evans, Stephen C. Levinson, Martin Haspelmath, Simon J. Greenhill, Quentin D. Atkinson, and Russell D. Gray. 2023b. [Grambank v1.0](#). Dataset.
- Richard Sproat, Bruno Cartoni, HyunJeong Choe, David Huynh, Linne Ha, Ravindran Rajakumar, and Evelyn Wenzel-Grondie. 2014. [A database for measuring linguistic information content](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 967–974, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Anirudh Srinivasan and Eunsol Choi. 2022. [TyDiP: A dataset for politeness classification in nine typologically diverse languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5723–5738, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Sabine Stoll and Balthasar Bickel. 2013. [Capturing diversity in language acquisition research](#). *Language typology and historical contingency*, pages 195–216.
- Clara Vania, Yova Kementchedjheva, Anders Søgaard, and Adam Lopez. 2019. [A systematic comparison of methods for low-resource dependency parsing on genuinely low-resource languages](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1105–1116, Hong Kong, China. Association for Computational Linguistics.
- Viveka Velupillai. 2012. *An Introduction to Linguistic Typology*, reprinted 2017 with corrections edition. John Benjamins Publishing Company, Amsterdam Philadelphia, PA.
- Ivan Vulić, Simon Baker, Edoardo Maria Ponti, Ulla Petti, Ira Leviant, Kelly Wing, Olga Majewska, Eden Bar, Matt Malone, Thierry Poibeau, Roi Reichart, and Anna Korhonen. 2020. [Multi-SimLex: A large-scale evaluation of multilingual and crosslingual lexical semantic similarity](#). *Computational Linguistics*, 46(4):847–897.
- Ekaterina Vylomova, Jennifer White, Elizabeth Salesky, Sabrina J. Mielke, Shijie Wu, Edoardo Maria Ponti, Rowan Hall Maudslay, Ran Zmigrod, Josef Valvoda, Svetlana Toldova, Francis Tyers, Elena Klyachko, Ilya Yegorov, Natalia Krizhanovsky, Paula Czarnowska, Irene Nikkarinen, Andrew Krizhanovsky, Tiago Pimentel, Lucas Torroba Hennigen, Christo Kirov, Garrett Nicolai, Adina Williams, Antonios Anastasopoulos, Hilaria Cruz, Eleanor Chodroff, Ryan Cotterell, Miikka Silfverberg, and Mans Hulden. 2020. [SIGMORPHON 2020 shared task 0: Typologically diverse morphological inflection](#). In *Proceedings of the 17th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 1–39, Online. Association for Computational Linguistics.

- Hongzhi Xu, Jordan Kodner, Mitchell Marcus, and Charles Yang. 2020. [Modeling morphological typology for unsupervised learning of language morphology](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6672–6681, Online. Association for Computational Linguistics.
- Ningyu Xu, Tao Gui, Ruotian Ma, Qi Zhang, Jingt-ing Ye, Menghan Zhang, and Xuanjing Huang. 2022. [Cross-linguistic syntactic difference in multilingual BERT: How good is it and how does it affect transfer?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8073–8092, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Aditya Yadavalli, Alekhya Yadavalli, and Vera Tobin. 2023. [SLABERT talk pretty one day: Modeling second language acquisition with BERT](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11763–11777, Toronto, Canada. Association for Computational Linguistics.
- Xinyu Zhang, Xueguang Ma, Peng Shi, and Jimmy Lin. 2021. [Mr. TyDi: A multi-lingual benchmark for dense retrieval](#). In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, pages 127–137, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xinyu Zhang, Nandan Thakur, Odunayo Ogundepo, Ehsan Kamalloo, David Alfonso-Hermelo, Xiaoguang Li, Qun Liu, Mehdi Rezagholizadeh, and Jimmy Lin. 2023. [MIRACL: A multilingual retrieval dataset covering 18 diverse languages](#). *Transactions of the Association for Computational Linguistics*, 11:1114–1131.

A Paper Retrieval Details

Venue (↓)	Date range	# papers	# search hits	# has_claim	# introduces_dataset	Date retrieved
AAAI	1980–2023	18,370	2	2	0	2024-01-02
ACL Anthology	1952–2023	92,006	188	105	36	2023-12-11
ICLR	2013–2023	5307	3	2	1	2024-01-02
ICML	2011–2023	8372	1	1	1	2024-01-02
IJCAI	2007–2023	8911	0	0	0	2024-01-03
NeurIPS	1987–2023	20,343	0	0	0	2023-12-23
Total	1952-2023	153,309	194	110	38	-

Table 3 – Detailed description of the dataset collected for annotation. Due to availability and quality reasons, not all proceedings of all years for all venues could be retrieved.

B Diversity Metrics

Mean Pairwise Distance (MPD) For a language sample L and feature vectors V , we retrieve v_l for $l \in L$. We then calculate the mean, pairwise Euclidean distances between all pairs in L :


$$MPD(L) = \frac{1}{|L|(|L| - 1)} \sum_{i,j \in L, i \neq j} \|v_i - v_j\|_2$$

where the feature vectors from V consists of the GEOGRAPHY, GENETIC or SYNTACTIC vectors in lang2vec (as shown per paper in Figure 5).


Feature Value Inclusion (FVI) For a language sample L and feature vectors V we look at every v_l for $l \in L$. We check if all possible, non-missing values (i.e., $x \notin \{\text{no_cov}, ?\}$, the values indicating ‘no coverage’ and ‘unsure’ in Grambank) for a particular feature are covered (\mathcal{V}_{F_i}). We do this for the set of all features in Grambank \mathcal{F} . Finally, we take the average over the total number of features:

$$FVI(L) = \frac{1}{|\mathcal{F}|} \sum_{i=1}^{|\mathcal{F}|} \frac{|\{v_l[F_i] : l \in L\}|}{|\mathcal{V}_{F_i}|}$$

C Long Tail of Included Languages



(a) Top portion of the long tail.



(b) Continued tail of included languages (not exhaustive).

Figure 8 – Long tail of included languages.

D ‘Typologically Diverse’ Datasets

Paper	Dataset name	L (↓)	MPSD	FVI	Task(s) or Topic
Vylomova et al. (2020)		90	0.76 [§]	0.92 [‡]	Morphological inflection
Henrichsen and Uneson (2012)	SMALLWorlds	53	0.65 [§]	0.94	ASR
FitzGerald et al. (2023)	MASSIVE	51	0.64 [§]	0.92	Slot-filling, Intent classification
Ruder et al. (2021)	XTREME-R	50	0.66 [§]	0.92 [‡]	Classification, Parsing, IR, QA
Kodner et al. (2022)		32 [†]	0.68 [§]	0.88 [‡]	Morphological inflection
Goldman et al. (2023)		26	0.64	0.89	Morphological inflection
Longpre et al. (2021)	MKQA	25	0.61	0.89	QA
Jiang et al. (2020)	X-FACTR	23	0.66 [§]	0.88	Knowledge retrieval
Bugliarello et al. (2022)	IGLUE	20	0.63	0.85	Visual QA, IR, RE, NLI
Dione et al. (2023)	MasakhaPOS	20	0.64 [§]	0.67	Part-of-speech tagging
Zhang et al. (2023)	MIRACL	18	0.64	0.85	IR
Palma Gomez et al. (2023)		16	0.62 [§]	0.84	Language learning, cloze test
Asai et al. (2022)		15	0.66	0.84	QA
Jancso et al. (2020)	ACQDIV	14	0.73 [§]	0.84	Language acquisition
Vulić et al. (2020)	Multi-SimLex	12	0.58	0.83	Semantic similarity
Eskander et al. (2020a)	MorphAGram	12	0.72 [§]	0.75	Morphological segmentation
Hennig et al. (2023)	MultiTACRED	12	0.58	0.84	Relation extraction
Ponti et al. (2020)	XCOPA	12	0.66	0.75 [‡]	Commonsense reasoning
Buechel et al. (2020)		12	0.58	0.72	Lexicon generation
Clark et al. (2020)	TyDiQA	11	0.67	0.76	QA
Zhang et al. (2021)	Mr. TyDi	11	0.67	0.76	IR
Shi et al. (2023)	MGSM	10	0.65	0.73	Arithmetic reasoning
Mott et al. (2020)		9	0.65	0.73	Morphological segmentation
Srinivasan and Choi (2022)	TyDiP	9	0.62	0.75	Politeness classification
Gupta et al. (2023)	Xtr-WikiQA & TyDi-AS2	8	0.64	0.69	QA, Answer sentence selection
Pfeiffer et al. (2022)	xGQA	8	0.58	0.66	Visual QA
Sproat et al. (2014)		8	0.58	0.75	Linguistic information content
Samir and Silfverberg (2023)		7	0.66	0.67	Morphological inflection
Ginn et al. (2023)		7	0.80 [§]	0.50	Interlinear glossing
Sarti et al. (2022)	DivEMT	6	0.64	0.69	MT post-editing
Glavaš et al. (2020)	XHate-999	6	0.64	0.49	Abusive language classification
Majewska et al. (2020)	SpAM	5	0.58	0.70	Verb analysis
Yadavalli et al. (2023)	SLABERT	5	0.61	0.66	Second language acquisition
Liu et al. (2021)	MaRVL	5	0.69	0.56	Visually grounded RE
Hung et al. (2022)	Multi2WOZ	4	0.61	0.58	Task-oriented dialogues
Bartelds et al. (2023)		4	– [§]	0.00	ASR
Sasaki et al. (2002)		2	0.84	0.39	Co-reference resolution

Table 4 – Papers that have a claim and introduce a dataset. **L** refers to the number of languages regarding the claim, with † indicating that one language does not have an ISO-639-3 code. **MPSD** is the Mean Pairwise Distance of syntactic URIEL features per language pair, with § indicating that at least one language has low (< 5%) syntactic feature coverage. **FVI** refers to the Grambank feature value inclusion of the languages in the dataset: are all possible, non-missing values for a feature in Grambank included using the language selection. Here, ‡ refers to at least one language not being covered in Grambank, these languages are excluded from the coverage metric. The **highest** and **lowest** values for **MPSD** and **FVI** are highlighted. Abbreviations: *QA* = Question Answering, *RE* = Reasoning, *NLI* = Natural Language Inference, *IR* = Information Retrieval, *ASR* = Automatic Speech Recognition, *MT* = Machine Translation.

E XTREME-R Results by Word Order

In Table 5 we can see that SVO is the most represented group for all tasks, with SOV coming in second. After that, the representation of features tapers off quickly, with (the rather rare) OSV, OVS, and VOS word orders having no coverage at all. We can see that the performance of models for SOV and VSO is consistently lower than SVO. This can be due to the difference in the number of languages included for the group, although we did not see this effect for the inflectional feature (Table 2).

Subtask	Model	Overall	By F	Δ	OSV	OVS	VOS	SVO	SOV	VSO	NDO	NA
MLQA [♠]	XLM-R-L	72.71 (7)	70.83 (7)	-1.88	(-)	(-)	(-)	75.22 (4)	70.80 (1)	67.00 (1)	70.30 (1)	(-)
	mBERT	61.30 (7)	57.14 (7)	-4.16	(-)	(-)	(-)	66.85 (4)	49.90 (1)	51.60 (1)	60.20 (1)	(-)
	mT5	75.59 (7)	74.06 (7)	-1.53	(-)	(-)	(-)	77.62 (4)	75.30 (1)	70.20 (1)	73.10 (1)	(-)
LArReQA [★]	XLM-R-L	40.75 (11)	39.31 (11)	-1.44	(-)	(-)	(-)	42.10 (6)	39.75 (2)	34.60 (1)	40.80 (2)	(-)
	mBERT	21.58 (11)	19.75 (11)	-1.83	(-)	(-)	(-)	24.10 (6)	15.10 (2)	17.00 (1)	22.80 (2)	(-)
TyDiQA [♠]	XLM-R-L	64.29 (9)	62.80 (9)	-1.49	(-)	(-)	(-)	67.26 (5)	58.55 (2)	62.60 (2)	(-)	(-)
	mBERT	58.36 (9)	56.91 (9)	-1.44	(-)	(-)	(-)	61.24 (5)	55.55 (2)	53.95 (2)	(-)	(-)
	mT5	81.94 (9)	81.45 (9)	-0.50	(-)	(-)	(-)	82.94 (5)	78.40 (2)	83.00 (2)	(-)	(-)
XQuAD [♠]	XLM-R-L	77.21 (11)	77.11 (11)	-0.10	(-)	(-)	(-)	76.70 (6)	76.55 (2)	74.40 (1)	80.80 (2)	(-)
	mBERT	65.05 (11)	63.55 (11)	-1.50	(-)	(-)	(-)	67.35 (6)	56.60 (2)	62.20 (1)	68.05 (2)	(-)
	mT5	81.54 (11)	81.04 (11)	-0.49	(-)	(-)	(-)	82.22 (6)	79.10 (2)	80.30 (1)	82.55 (2)	(-)
XNLI [♠]	XLM-R	79.24 (15)	78.57 (15)	-0.67	(-)	(-)	(-)	80.31 (9)	75.10 (3)	77.20 (1)	81.65 (2)	(-)
	mBERT	66.51 (15)	65.79 (15)	-0.72	(-)	(-)	(-)	68.19 (9)	60.03 (3)	66.00 (1)	68.95 (2)	(-)
	mT5	84.85 (15)	84.71 (15)	-0.14	(-)	(-)	(-)	85.39 (9)	81.83 (3)	84.50 (1)	87.10 (2)	(-)
Mewsl-X [★]	XLM-R-L	45.75 (11)	45.66 (11)	-0.09	(-)	(-)	(-)	53.16 (5)	35.98 (4)	28.70 (1)	64.80 (1)	(-)
	mBERT	38.58 (11)	37.88 (11)	-0.71	(-)	(-)	(-)	47.28 (5)	27.93 (4)	15.30 (1)	61.00 (1)	(-)
Tatoeba [♠]	XLM-R	77.29 (41)	72.82 (38)	-4.46	(-)	(-)	(-)	81.42 (18)	75.74 (14)	64.55 (2)	86.57 (4)	55.83 (3)
	mBERT	43.33 (41)	39.66 (38)	-3.67	(-)	(-)	(-)	52.57 (18)	33.04 (14)	25.10 (2)	54.78 (4)	32.83 (3)
XCOFA [♠]	XLM-R	69.22 (11)	66.16 (9)	-3.06	(-)	(-)	(-)	72.89 (7)	72.90 (2)	(-)	(-)	52.70 (2)
	mBERT	56.05 (11)	55.17 (9)	-0.88	(-)	(-)	(-)	57.11 (7)	55.50 (2)	(-)	(-)	52.90 (2)
	mT5	74.89 (11)	72.62 (9)	-2.27	(-)	(-)	(-)	77.61 (7)	77.00 (2)	(-)	(-)	63.25 (2)
WikiANN-NER [♠]	XLM-R-L	64.43 (48)	65.10 (43)	+0.67	(-)	(-)	(-)	66.80 (20)	59.74 (17)	57.95 (2)	79.70 (4)	61.30 (5)
	mBERT	62.68 (48)	63.98 (43)	+1.30	(-)	(-)	(-)	67.34 (20)	54.26 (17)	58.80 (2)	75.92 (4)	63.58 (5)
UD-POS [♠]	XLM-R-L	74.96 (38)	77.45 (36)	+2.50	(-)	(-)	(-)	74.59 (20)	69.65 (10)	71.55 (2)	86.97 (4)	84.50 (2)
	mBERT	70.90 (38)	71.60 (36)	+0.69	(-)	(-)	(-)	72.74 (20)	62.82 (10)	61.25 (2)	83.22 (4)	77.95 (2)

Table 5 – XTREME-R task and model results, grouped by word order (WALS 81A). **Overall** refers to the average over all languages. **By Feature** is the average of the WALS feature averages, excluding languages for which there is no coverage. The delta shows the difference of **Overall** and **By F**. Tasks that have coverage for all their included languages are on top, those partially covered are in the bottom portion. The word orders are listed in the following columns, with *NDO* being *No Dominant Order* and *NA* referring to *Not Available* (in WALS). The **highest** and the **lowest** number of languages per grouping are highlighted. The *(italicized)* number refers to the number of languages in a particular subset. Metrics: \blacklozenge = Accuracy, \blackstar = mAP@20, \blackspade = F1.

F XTREME-R Grambank Feature Value Inclusion



Figure 9 – PCA Grambank feature value inclusion per XTREME-R subtask. Where applicable, we cite the original dataset that XTREME-R includes or used as a starting point to expand upon.