

A Principled Framework for Evaluating on Typologically Diverse Languages

Esther Ploeger¹, Wessel Poelman², Andreas Holck Høeg-Petersen¹, Anders Schlichtkrull¹, Miryam de Lhoneux², and Johannes Bjerva¹

¹Department of Computer Science, Aalborg University

espl@cs.aau.dk, ahhp@cs.aau.dk, andsch@cs.aau.dk, jbjerva@cs.aau.dk

²Department of Computer Science, KU Leuven

wessel.poelman@kuleuven.be, miryam.delhoneux@kuleuven.be

Beyond individual languages, multilingual natural language processing (NLP) research increasingly aims to develop models that perform well across languages generally. However, evaluating these systems on all the world's languages is practically infeasible. To attain generalizability, representative language sampling is essential. Previous work argues that generalizable multilingual evaluation sets should contain languages with diverse typological properties. However, "typologically diverse" language samples have been found to vary considerably in this regard, and popular sampling methods are flawed and inconsistent. We present a language sampling framework for selecting highly typologically diverse languages given a sampling frame, informed by language typology. We compare sampling methods with a range of metrics and find that our systematic methods consistently retrieve more typologically diverse language selections than previous methods in NLP. Moreover, we provide evidence that this affects generalizability in multilingual model evaluation, emphasizing the importance of diverse language sampling in NLP evaluation.

1. Introduction

Data-driven approaches to language technology have shifted the realm of possibility in multilingual NLP. Distributed word representations (Mikolov et al. 2013) have lifted the reliance on language-specific hand-crafted rules. This is leveraged by pre-training language models on multiple languages simultaneously, such as multilingual BERT (Devlin et al. 2019) and XLM-R (Conneau et al. 2020), increasing performance through transfer learning. More recently, even English-centric large language models (LLMs) are claimed to "possess multilingual capabilities that surpass our expectations" (Yuan et al. 2023), in the case of LLama (Touvron et al. 2023), and "effectively transfer learned knowledge across different languages" (Zhang et al. 2023), in the case of GPT3.5.

A shared factor behind the success of these models is their reliance on large volumes of textual data. These data-driven approaches are claimed to be language-agnostic, as

Action Editor: Michael Roth. Submission received: 5 July 2024; revised version received: 5 September 2025; accepted for publication: 21 September 2025.

<https://doi.org/10.1162/COLLa.577>

they are, in principle, applicable to any language given enough training data. However, language-agnostic systems are not language-independent (Bender 2009, 2011). Current algorithms are designed with an English-centric perspective in mind, while characteristics (such as low morphological complexity) of the English language cannot be assumed to transfer to other languages. To the best of our knowledge, there is currently no systematic investigation of how the properties of languages included in the evaluation influence the performance estimations of multilingual language models. It remains unclear how well these models perform on diverse languages systematically sampled.

To assess whether a language model performs well *across* languages, ideally, one would evaluate it on all languages in the world. However, collecting high-quality data on such a scale is not feasible. Therefore, multilingual models are evaluated on a sample of the world's languages. To increase generalizability, such a language sample should be diverse, with varying characteristics and properties. However, existing approaches to this sampling process have not been optimal. Previous work has established that there is no clear terminology or methodological consensus for what constitutes "typological diversity" (Ploeger et al. 2024). Currently, many approaches in NLP use phylogenetic heuristics to ensure typological diversity. In this article, we show that this approximation of typological diversity through language phylogeny has severe shortcomings, and we provide a framework for systematically selecting languages based on typological distance measures. Our framework implements two sampling methods that are widely established within the field of linguistic typology. We demonstrate how our framework can be used for diverse typological language sampling, how it can help guide dataset expansions, and how it has use beyond typology.

Contributions. (i) We establish that language phylogeny is limited when it comes to assessing typological diversity; (ii) we provide a method for quantifying the typological distance between pairs of languages; (iii) we provide a systematic framework for selecting typologically diverse languages for multilingual evaluation scenarios; (iv) we introduce four measures of typological diversity; (v) we show that typological diversity matters in downstream evaluation; and (vi) we provide a Python package that facilitates typologically diverse language selection, which is publicly available in the following repository: <https://github.com/esther2000/typdiv-sampling>.

2. Background

Representative language sampling has long been a central methodological issue in the field of linguistic typology. In order to assess its applicability to sampling in NLP, we discuss this line of research in more detail.

Linguistic typology can be described as the study of structural similarities and differences across the world's languages (Kashyap 2019). Within linguistic typology, an important research direction is the investigation of general patterns for these similarities and differences. For example, Greenberg (1963) finds that in languages with verb-object word order, adpositions tend to be placed before their objects. For drawing such general conclusions about human languages as a whole, testing linguistic hypotheses on only a handful of related languages is insufficient (Rijkhoff et al. 1993; Guzmán Naranjo and Becker 2022). Instead, generalizable conclusions in typology require adequate sampling strategies: Findings should be supported by evidence across a diverse range of languages. Thus, language sampling has received considerable attention in typological research. Given the recent advances in language technology beyond English, language

sampling has become increasingly relevant for multilingual NLP. Drawing generalizable conclusions about the performance of multilingual models requires tests across diverse languages. Moreover, gaining insight into the skews of evaluation language sets may help to identify weaknesses of current applications. In this section, we discuss relevant sampling strategies and terminology from the field of linguistic typology and how they relate to language sampling in NLP, despite differences in objectives.

Universe, Frame and Sample. Bell (1978) introduced central notions for language sampling in typology. To establish a common ground for approaching language sampling more systematically in NLP, we relate these terms to NLP. Firstly, the sampling UNIVERSE is the class of objects under study. In typological studies, this could for instance be the set of all possible human languages. In NLP, the sampling universe often corresponds to all existing natural languages that a certain language technology is claimed to generalize to. The sampling FRAME provides access to the sampling universe. It is the concrete set of languages one can actually sample from. Sjöberg (2023) further describes the distinction between a **catalogue** frame, which is the set of languages we know to exist, and a **corpus** frame, which is the set of languages for which we have relevant research materials. For our purposes, we will treat the catalogue frame as all languages we have typological information for. We will treat the corpus frame as the set of languages that we have available datasets for. Given the data requirements of most techniques in NLP, the corpus frame is of vital importance. Lastly, the SAMPLE is the set of languages that one draws from the frame, with the aim to reflect the properties of the sampling universe (Figure 1). In NLP, this can correspond to the concrete set of languages that one tests or evaluates an application on.

Sampling Methods. For extracting a sample from the frame, there are multiple popular types of methods in linguistic typology. Rijkhoff and Bakker (1998) divide sampling methods into three categories: RANDOM, PROBABILITY, and VARIETY sampling. Each of these sampling methods is relevant for different research questions. Random sampling entails selecting languages without any criteria. Without “stratification”, the grouping of the frame before sampling, it is possible that resulting samples are in large part made

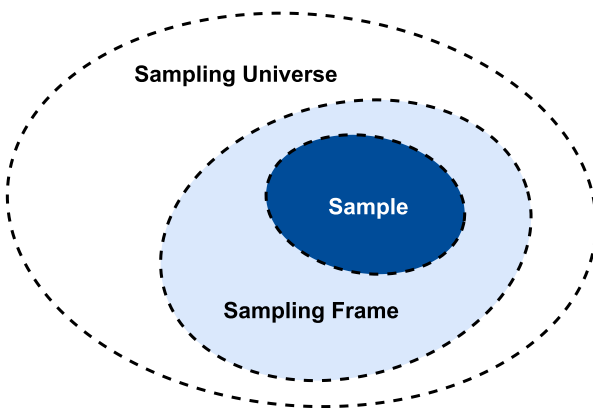


Figure 1
A diagram showing the relationship between the sampling universe, sampling frame, and the sample. The objective of diverse language sampling is to draw a sample from the sampling frame that optimally reflects a larger population (frame or universe).

up of similar languages. As such, these samples could be skewed towards languages from certain phylogenetic or geographical groups. This is especially likely if the sampling frame contains a skew. Such samples, given they are large enough, are commonly used to look into the occurrence frequency of some linguistic phenomenon, but not necessarily for generalizable conclusions about language. In probability sampling, one aims for a sample that contains independent languages. Ideally, the resulting sample should be free of bias. For instance, it should not contain a skew towards one or a few language families. Variety sampling aims at sampling languages such that the linguistic diversity of the world's languages is captured as much as possible. This is because "exceptional types test the rule" (Perkins 1988). It should be noted, however, that rather than creating statistical models of all languages of the world, sampling in NLP is rather concerned with testing multilingual model robustness across languages. In the context of NLP, Ponti et al. (2020) write that choosing a variety sample tests the robustness of a language model to unseen typological features, as it includes outliers. These three sampling methods each come with different implications in terms of sample size. Because of the lack of stratification, a random sample "must be relatively large in size to be able to produce reliable results" (Rijkhoff and Bakker 1998). For variety sampling, it is also often the case that the more languages are included, the better. This is because including more languages means that outliers and uncommon properties are more likely to be captured (Miestamo, Bakker, and Arppe 2016). However, for probability sampling, there is a trade-off between independence and coverage. The more languages one includes, the more difficult it becomes to preserve independence between these languages. To illustrate: If the number of languages one samples is larger than the number of language families in the frame, then multiple languages from the same family have to be sampled.

Language Sampling in NLP. Bender (2009, 2011) and Pikuliak and Simko (2022) extensively discussed the need for diverse evaluation language selection, and the potential for linguistic typology to facilitate this. However, to date, there are only a handful of works in NLP that aimed to apply these suggestions. In fact, Ploeger et al. (2024) ascertained that typologically diverse sampling methods in NLP are mostly flawed. Firstly, if any stratification criterion is mentioned in NLP, this is usually based on phylogeny (e.g., Yadavalli, Yadavalli, and Tobin 2023; Ács, Kádár, and Kornai 2021; Xu et al. 2020). To the best of our knowledge, no approach in NLP uses a tree structure to apply genealogical stratification. Instead, sampling from different families is common. To illustrate: Majewska et al. (2020) "sampled languages from 5 different language families to ensure typological diversity". However, there is no evidence that phylogenetic relations directly imply typological relations (Dahl 2008). We further explore this in Section 3. Secondly, most sampling stratification is applied post-hoc. Commonly, a set of languages is selected, and only then is it described in terms of diverse language families and sometimes typological diversity. For example, the typology index by Ponti et al. (2020) (see Section 6.2) is only applied after language selection. As such, it does not promote principled language selection. Our work differs from these approaches, in that we perform informed *selection*, instead of only informed *analysis*.

3. Why Is Phylogeny Insufficient?

In Section 2, we mentioned how previous work in NLP typically approximates typological diversity through phylogenetic groupings. More specifically, a major share of approaches claim to approximate typological diversity by randomly selecting languages

from different language families or genera. However, it is not evident that phylogenetic relationships directly imply typological similarities (Georgi, Xia, and Lewis 2010). Furthermore, relying on this stratification criterion can give vastly inconsistent samples (and thus different results), given that within-family selection is random. In this section, we critically assess the shortcomings of phylogeny as a proxy for typologically diverse language sampling in NLP.

Theoretical Arguments. Phylogenetic groupings such as language families and genera are commonly described as strictly exclusive groups. For example, German, Dutch, and Hindi are all in the Indo-European language family. In reality, language similarity is much more gradient. Intuitively, it may be clear for some that German and Dutch are much more similar than Dutch and Hindi. It is therefore not surprising that the strict boundaries between families are not necessarily agreed upon. For example, Dixon (1997) writes “about 1,000 languages have been grouped together in a putative ‘Niger-Congo family’. [...] One searches in vain for proof of this ‘genetic relationship.’” Furthermore, Dahl (2008) writes that “[genealogically] related languages that are no longer in contact with each other can in a few thousand years develop typological profiles that are no longer indicative of a common origin.” The imposed strict distinction between phylogenetic language groups also causes other issues. For instance, pidgins and Creoles are not easily placed into a family tree, which means that sometimes they are excluded from sampling methods. For instance, Sjöberg (2023) writes: “Finally, pidgins and Creoles are also excluded from the frame, due to the difficulty of deciding where in and in what family tree to place them.” Furthermore, strictly divided language families come in different sizes. The Niger-Congo language family includes more than 1,000 languages (Hammarström et al. 2024), while isolates such as Basque constitute their entire “family”. This further complicates fair sampling. Finally, if the number of languages one wishes to sample is smaller than the number of phylogenetic groupings, there is a random selection on a group-level as well.

Empirical Arguments. Beyond theoretical reasons, it is unclear to what extent phylogenetic groupings actually imply typological diversity. Here, we empirically assess the extent to which typological similarity overlaps with language families. Grambank (Skirgård et al. 2023) contains genealogical and typological information for 2,467 languages. We use this data to calculate, for each language family in Grambank, the pairwise average of overlapping feature values in that family. In Figure 2, we show this overlap per language family, as well as the number of languages in each family. If families were coherent typological groups, we would expect to see high overlap throughout—we observe variety in overlap, with most averages below 0.6. Additionally, we retrieve the closest language for each language in Grambank, in terms of feature value overlap. We find that the closest language is in another family in 32.42% of the cases. Thus, sampling from distinct language families does not directly imply that the sampled languages are typologically distant.

All in all, we conclude that sampling with phylogenetic stratification is not ideal for NLP purposes. In linguistic typology, sampling using a proxy is often necessary. Directly using typological values instead is to be avoided, because then using those same variables for sampling introduces circularity. However, in NLP, the variable under investigation is commonly something different, such as the performance of multilingual models. Thus, there is no circularity with sampling directly from typological features. This avoids many of the issues with strict group sampling: We can take into account granularity, include languages that are typically difficult to place in one of those strict




Figure 2
 Number of languages and average within-family feature value overlap for the 25 largest language families in Grambank.

groups, we can control the group size, and are not affected by randomness in sampling. Additionally, by sampling through typological features directly, we gain immediate insight into the typological diversity of the sample. In Section 5, we present a method that does exactly this.

4. Related Work

We argue that typologically generalizable evaluation in multilingual NLP should be conducted on the basis of a priori sampling with typological features. To the best of our knowledge, this is only implemented in two previous works. Dahl (2008) sampled a subset of the World Atlas of Language Structures (WALS; Dryer and Haspelmath 2013) by removing one language from each pair that is above a certain typological similarity threshold. The language that is removed is the one with the least coverage in WALS, which exacerbates bibliographical bias (Bakker 2010). Stoll and Bickel (2013) introduce a sampling method based on fuzzy clustering (Kaufman and Rousseeuw 2009). They divide languages into k clusters based on twelve typological features, which the authors manually coded from various sources. Then, for each cluster, they sample the language with the highest membership coefficient.

Our work differs from these approaches considerably. Firstly, our method is designed to handle data-inequality. This means that it does not assume that all typological features are described for every language, contrary to previous work. As complete typological coverage is rare, this renders our method applicable to a far broader range of languages and typological features. Also, we present a flexible framework, where individual features can be left out, if desired. Importantly, our framework accommodates both probability sampling and variety sampling from linguistic typology. This is important, because different sampling methods may be used to answer different research questions.

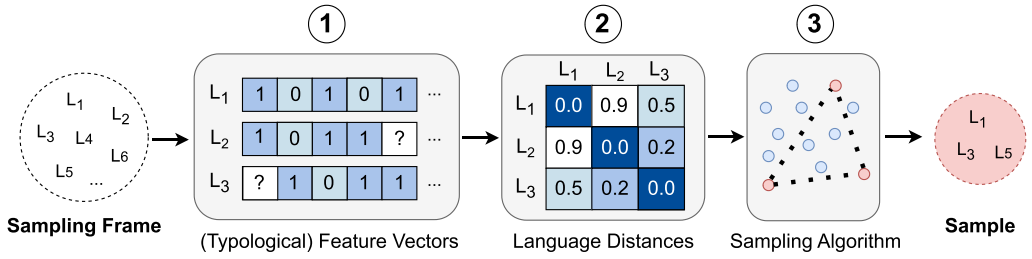


Figure 3 Language sampling by measuring distances of typological feature vectors.

5. A Principled Language Sampling Framework

The task of language sampling consists of selecting a set of languages (sample) from a larger set of languages (frame). There are two main methods for performing typologically diverse sampling. For variety sampling, the languages in the sample should be maximally diverse. For probability sampling, these languages should be maximally independent. Instead of approximating typological diversity through phylogenetic relationships, we measure typological distance between languages based on typological properties. Our approach consists of three steps:

1. Retrieve typological information per language (Section 5.1)
2. Calculate pairwise distances between languages (Section 5.2)
3. Sample languages using an algorithm that calculates a set of typologically distant languages (Section 5.3)

In Figure 3, we schematically represent our complete sampling pipeline. In the next subsections, we discuss each of the separate steps in more detail.

5.1 Typological Feature Vectors

We use the Grambank database (Skirgård et al. 2023) for retrieving typological characteristics of the languages in a sampling frame. Grambank v1.0, which we use throughout this article, contains typological information for 2,467 language varieties, for 195 grammatical features. Of these features, 189 are expressed as binary statements (e.g., *Are there definite or specific articles?*; GB020). These features can take values 0, 1, or ?, where the latter denotes unclear or unknown features. Coverage is incomplete: For some language varieties, some features are not described, indicated by no_cov. Six word-order features are multi-value. For instance, feature GB024 (*What is the order of numeral and noun in the NP?*) can take the values Num-N, N-Num, both, and ?.

Our framework is compatible with any kind of information about languages. The framework can also be used starting from step 3 when there are already pairwise

distances available. These language distances are not required to be based on typology.¹ In this work, we specifically use Grambank, for multiple reasons. Firstly, Grambank was developed with computational applications in mind, taking specific care to avoid logical dependencies between feature values (Haynie et al. 2023). Logical dependencies are introduced when the value of one feature implies another. These are problematic for calculating (language) distances, as features that directly imply each other are then weighted more than features that are not logically implied by others.² Additionally, since our work focuses on text processing specifically, the morphosyntactic domain of Grambank is particularly relevant. For our main experiments, we therefore treat Grambank as our catalogue frame. Still, if one aims to conduct a phonological study, one could use phonological features with our framework instead. It should be noted that the representativeness of this sampling approach heavily depends on the input features. For example, using input features that are only available for a few languages cannot retrieve representative samples more generally. In addition, statistical dependencies between features can bias the sampling results, making it essential to relate the input features to the use case at hand.

5.2 Language Distance Calculation

Let a sampling frame be a set \mathcal{L} of languages. For each language $l \in \mathcal{L}$ we extract a typological vector $V(l)$ with d dimensions, which consists of all Grambank feature values for l . We binarize the six multi-value word order features, as suggested by Haynie et al. (2023). This leaves us with $d = 209$ features which for each language will have values 0, 1, ?, or `no_cov`, the latter representing a feature not covered. We treat both features without coverage and features with a value of ? as explicitly missing values, which we both treat as NaN. For a vector v and integer i , let v_i be the i th feature of v . For each pairwise combination of languages $l, l' \in \mathcal{L}$, we then calculate the Euclidean distance in the presence of missing values, as defined by Dixon (1979) and implemented by Pedregosa et al. (2011). This distance calculation is defined in Equation (1).

$$\text{dist}(l, l') = \sqrt{w(l, l') \cdot \sum_{f \in s(l, l')} (V(l)_f - V(l')_f)^2} \quad (1)$$

where $s(l, l')$ is the set of features that are covered in both l and l'

$$s(l, l') = \{f \in \{1 \dots d\} \mid V(l)_f \neq \text{NaN} \text{ and } V(l')_f \neq \text{NaN}\} \quad (2)$$

and where w is calculated by dividing the total number of data points by the number of present data points:

$$w(l, l') = \frac{d}{|s(l, l')|} \quad (3)$$

¹ We show an application of this in Section 8.2.

² It should be noted that a lack of *logical* dependencies does not exclude statistical dependencies, as some feature values may co-occur.

This provides us with the distance between all pairs of language varieties in Grambank. We use these distances in our algorithms for calculating sets of typologically distant languages (Section 5.3). Our framework provides additional flexibility by supporting a number of processing steps:

- **Normalization:** For better readability of the pairwise distances, min-max normalization can be applied, retrieving distances in the range $[0, 1]$. This does not affect the sampling process or results.
- **Binarization:** For using Grambank’s multistate values in the same format as the binary features, we incorporate a binarization option. We follow the authors of Grambank by dividing each multistate value into two binary features, which are not logically dependent.³
- **Language cropping:** In order to mitigate influences of languages with very limited feature coverage on the sampling results, we provide the option of removing languages with a defined percentage of missing data. Following Skirgård et al. (2023), we use a threshold of $>25\%$ missing data for language cropping in the rest of this work.
- **Removing macro-languages:** Grambank contains a number of macro-languages (e.g., Central pacific linkage, Oceanic), which may not be relevant to one’s case study. Our framework facilitates filtering these out, informed by the number of child languages in Glottolog.
- **Feature sub-selection:** For case studies that only comprise a subset of morphosyntactic features, our framework supports the selective inclusion of features.

For the demonstrated applications of our framework in this work, we apply normalization, binarization, and language cropping and remove macro-languages.

5.3 Sampling Algorithms

Given these typological distances between pairs of languages, there are multiple ways to do typologically diverse language selection. In the linguistic typology literature, different sampling methods are used for answering different kinds of research questions. Inspired by this, our framework explicitly accommodates both variety sampling and probability sampling methods. While we specifically focus on typology in this section, these methods can be used with any type of information, as mentioned in Section 5.1.

5.3.1 MaxSum Diversity. For our typologically informed *variety* sampling, we treat the problem as a MAXIMUM DIVERSITY PROBLEM (MDP).⁴ This entails finding a size k set of points where the sum of distances between all points in the set is maximal (*MaxSum*). Approaching the problem as such optimizes for sampling for large *total* distance, capturing outliers (see Figure 4), which is the objective of variety sampling.

³ https://github.com/grambank/grambank/blob/master/docs/Grambank_most_updated_sheet.tsv.

⁴ Marti et al. (2013) list a number of alternative names used for the problem: maximum dispersion, max-avg dispersion, p -dispersion, p -dispersion-sum, edge-weighted clique, remote clique, maximum edge-weighted subgraph, dense k -subgraph, p -defense, p -defence-sum, and equitable dispersion.




Figure 4

A visualization of both sampling algorithms, with the MaxSum objective on the left and MaxMin on the right over a normal distribution. The red triangles represent the sample selected by the respective algorithm, the blue dots the remaining languages in frame. The distance here is the Euclidean distance in a 2D plane.

Kuo, Glover, and Dhir (1993) showed that the “clique problem”, finding subsets of vertices in a graph that are all adjacent, can be reduced to MDP. Since the clique problem is NP-hard, then so must MDP be. Famously, there is no known efficient algorithm for finding optimal solutions to NP-hard problems. We also experienced that a brute force algorithm did not terminate when run on the dataset. Martí et al. (2013) give an overview of heuristics for MDP and conclude that even simple heuristics give good solutions. We implement one such simple heuristic. Given pairwise language distances, we first find the language that is most distant from all others. For this, we take the language where the sum of distances with all other languages is largest. The motivation behind this is that selecting the language with the most “unusual” typological properties (or combinations thereof) already ensures some outlier. Next, we add the language that is furthest from the first language to the sample. We then sum the distances from these two languages to every other language. The language that has the largest summed distance is added to the selection. We repeat this process until the desired size of the sample is reached. The full procedure is described in Algorithm 1.

5.3.2 MaxMin Diversity. In sampling with the MaxSum objective, the total typological distance between languages is maximized. Individual outliers within the sample may be close (see Figure 4), which may introduce a skew in a sample.⁵ However, for *probability* sampling, one would want to preserve the independence between languages. To this end, we approach the diversity problem as a MAXMIN DIVERSITY PROBLEM, which is a second popular diversity sampling objective (Parreño, Álvarez-Valdés, and Martí 2021). Instead of optimizing for the maximum total distance, we optimize for maximizing the smallest distance between any two points in the set. Thus, we select languages such that the closest two are maximally typologically distant. This problem is NP hard as

⁵ The distribution of language types can be visualized as distributed normally around the center; see Dryer (1998).

Input: k : number of languages to sample, \mathcal{L} : sampling frame, $dist$: function giving the pairwise distance between languages in \mathcal{L} (e.g., Equation (1))

Algorithm 1 MaxSum Sampling

```

1:  $l \leftarrow \arg \max_{l \in \mathcal{L}} \sum_{l' \in \mathcal{L}} dist(l, l')$ 
2:  $L \leftarrow \{l\}$ 
3: while  $|L| < k$  do
4:    $L \leftarrow L \cup$ 
      $\{\arg \max_{l \in \mathcal{L} \setminus L} \sum_{l' \in L} dist(l, l')\}$ 
5: end while
6: return  $L$ 

```

Algorithm 2 MaxMin Sampling

```

1:  $l \leftarrow \arg \max_{l \in \mathcal{L}} \sum_{l' \in \mathcal{L}} dist(l, l')$ 
2:  $l' \leftarrow \arg \max_{l' \in \mathcal{L}} dist(l, l')$ 
3:  $L \leftarrow \{l, l'\}$ 
4: while  $|L| < k$  do
5:    $L \leftarrow L \cup$ 
      $\{\arg \max_{l \in \mathcal{L} \setminus L} (\min_{l' \in L} dist(l, l'))\}$ 
6: end while
7: return  $L$ 

```

well, which means we again need to rely on heuristics. Similar to MaxSum diversity, we implement a simple heuristic. As in the MDP approach, we first select the language that is most distant from all others. Then, we again take the language furthest from that language. Then until the desired sample size (k) is reached, we add the next language to the sample that retrieves the highest minimum distance to the already selected points. This last step is the key difference between the objectives of both algorithms. The full process is described in Algorithm 2.

Within linguistic typology, variety sampling is used to test a hypothesis “against the extremes”, while probability sampling is used for research questions that require more even language distributions. In Section 7, we further contrast these methods for downstream use cases in NLP, showing that MaxSum sampling is especially useful for scenarios in which *a worst-case performance estimate is desired*.

6. Typological Diversity Evaluation

Our sampling methods enable a priori language sampling for typological diversity. In Section 8 we look into how our algorithm can be used practically. Here, we first verify whether our sampling methods actually retrieve more diverse language samples than previous methods used in NLP. To this end we compare our sampling algorithms to approaches from previous work in multilingual NLP with four metrics. Additionally, we address the issue of how many languages one would need to include for a representative sample.

6.1 Baselines

Random. As mentioned in Section 2, language selection in NLP is mostly unprincipled, where evaluation languages are sampled without stratification methods. This resembles random sampling in linguistic typology. In typology, random samples are mostly used for gaining insight into occurrence frequencies, but not for drawing generalizable conclusions. For comparability, we include a baseline which randomly samples from the sampling frame (Grambank).




Figure 5 Schematic overview of diversity metric calculations, where “for each pair” means all combinations of languages in the set.

Convenience. In practice, random sampling in NLP is not truly random, because data availability plays a large role. Therefore we select the languages that are most commonly used in previous work in NLP that claim to have typologically diverse language selections. Ploeger et al. (2024) annotated the language selections in papers containing such claims. For our baseline, we sort these languages based on occurrence frequency, and sample the first k . In case of a tie in terms of occurrence frequency, selection is random.

Phylogenetic. The most popular stratification in NLP, if any, is sampling based on language families. This method can be motivated from the perspective of linguistic typology, where phylogenetic stratification is popular. Contrary to linguistic typology, advanced sampling methods from language trees are not relevant to NLP (Section 2). We implement phylogenetic stratification by sampling uniformly from groupings on two levels: language families⁶ and genera.⁷ If the requested sample size is bigger than the number of groupings, we uniformly sample again from the groupings until we reach the requested size. Note that this is already more principled than most sampling approaches in NLP. The convenience baseline is indirectly also often influenced or motivated by phylogeny, albeit in a less principled manner. As such, our phylogenetic baseline is a “best case scenario” for principled, typologically diverse language sampling using phylogenetic data.

6.2 Metrics

Measuring the typological diversity of a language set can be done in a number of ways. We formulate and compare four diversity metrics here, incorporating both previous work and new ideas. We indicate “higher is better” metrics with an up arrow (\uparrow) and “lower is better” with a down arrow (\downarrow). Figure 5 provides a schematic representation of the metrics.

Mean Pairwise Distance (MPD). The mean pairwise distance (MPD) of a language set measures the Euclidean distance between all combinations of languages in the sample

⁶ As specified in Glottolog v4.4: 215 families.

⁷ As specified in WALS v2020.3: 612 genera.

(Ploeger et al. 2024). Because these distances are used in our sampling algorithms directly, this measure serves as a sanity check to verify whether the distances that the algorithm is based on are actually increased. Pairwise comparisons can be motivated from a typological perspective (Wichmann and Holman 2010), and taking the average means that the results can be compared across language samples of different sizes. MPD can be formalized as follows:

$$MPD(L) = \frac{1}{|L|(|L| - 1)} \sum_{l, l' \in L, l \neq l'} dist(l, l') \quad (4)$$

Here $dist(l, l')$ denotes the Euclidean distance as defined in Equation (1).

Feature Value Overlap (FVO). Distances alone do not directly describe the disparity of our features. This is motivated from the perspective of linguistic typology; Dahl (2008) describes measuring language similarity as: “How large a proportion of the features that are defined for both members of a language pair have different values?”. To this end we calculate the feature value overlap (FVO), which is the average of the percentages of features that overlap between any pair of languages in the combinations of a language set. Since Grambank contains binarized feature values (as outlined in Section 5.2), calculating such an overlap is appropriate.⁸ We report the average over all combinations. Feature value overlap can be formalized as follows:

$$FVO(L) = \frac{1}{|L|(|L| - 1)} \sum_{l, l' \in L, l \neq l'} \frac{|\{f \in \{1 \dots d\} \mid V(l)_f = V(l')_f\}|}{d} \quad (5)$$

Feature Value Inclusion (FVI). The previous metrics do not measure the extent to which individual typological properties are covered. This is especially relevant for variety sampling, where rare typological features should be included. Miestamo, Bakker, and Arppe (2016) defined the measure of *saturation* of a typological feature as “the proportion of values, out of the maximum number of possible values, found in the sample for that feature”. As the name suggests, this metric is expected to saturate for larger language selections. Therefore, it is especially suitable for comparing smaller language samples. Because we deal with binary features only, we calculate the feature value inclusion (FVI) per feature as the percentage of languages that include the feature. We report the average over all features. FVI can be formalized as follows:

$$FVI(L) = \frac{1}{d} \sum_{f=1}^d \frac{|\{V(l)_f \mid V(l)_f \neq \text{NaN and } l \in L\}|}{2} \quad (6)$$

Entropy (H). Similar to the diversity index reported in Ponti et al. (2020), we report the entropy of the feature values that occur in a language sample. This gives us insight into the spread within features. For example, if a sample L_1 contains $[1, 1, 1, 1, 0]$ for a given feature f , FVI is maximal, but the spread is low. The entropy for f is lower than that of a more diverse sample L_2 with values $[1, 1, 0, 0, 0]$. We take the average of this metric over all features. A difference with previous work is that Ponti et al. (2020) based

⁸ Special care has to be taken to calculate this metric when using multistate features values.

their entropy calculations on 103 unnamed typological features from URIEL (Littell et al. 2017). This is not reproducible, as they do not report which typological features were used. Furthermore, as URIEL contains logical dependencies, the included typological features are not necessarily weighted equally. The entropy of a set of languages is the average entropy over all features:

$$H(L) = \frac{1}{d} \sum_{f=1}^d H(f) \quad (7)$$

where the entropy of a feature is

$$H(f) = - \sum_{i \in \{0,1\}} p(f,i) \cdot \log_2 p(f,i) \quad (8)$$

and the probability p is calculated as

$$p(f,i) = \frac{|\{l \in L \mid V(l)_f = i\}|}{|\{l \in L \mid V(l)_f \neq \text{NaN}\}|}. \quad (9)$$

6.3 Results and Discussion

We compare our sampling methods against baselines based on previously used methods in NLP, as introduced in Section 6.1.

Figure 6 shows that our methods consistently retrieve more diverse samples than all baselines. For the pairwise metrics (MDP, FVO), this difference is especially large for smaller samples; it is easier to avoid overlap when sampling fewer languages. This is in line with earlier findings from typology, which described the trade-off between coverage and independence (Miestamo, Bakker, and Arppe 2016). While the FVI difference across methods is small for large sample sizes, we find that our methods retrieve a higher inclusion of feature values for small samples (<20) than the baselines. This is especially relevant for NLP, where multilingual hypotheses are commonly tested on a small set of typologically diverse languages (median = 11; Ploeger et al. 2024). Lastly, we observe that MaxSum and MaxMin consistently retrieve samples with higher feature entropy than other methods.

These results also showcase how our methods can be used to inform the choice of sample size in the design of a study. For these particular corpus frames, it can be argued that a sample size of 20 languages adequately covers the feature space. For instance, FVI flattens after $k > 20$. This could be a justification for making claims about the generalizability of a certain phenomenon captured in a sample.⁹

Next, we compare the methods in more detail, for a specific value of k . We zoom in on $k = 20$, since the metrics tend to flatten off for all methods around that value, as seen in Figure 6. We measure the MPD, FVO, FVI, and H for the samples that all our sampling approaches retrieve. We run the non-deterministic methods 10 times. The results (Table 1) show that both MaxSum and MaxMin sampling retrieve considerably better results than all baselines.

⁹ While the number of languages one can sample from is a characteristic of the method, we find similar results when taking the intersection of all methods' sampling frames: see Appendix A.

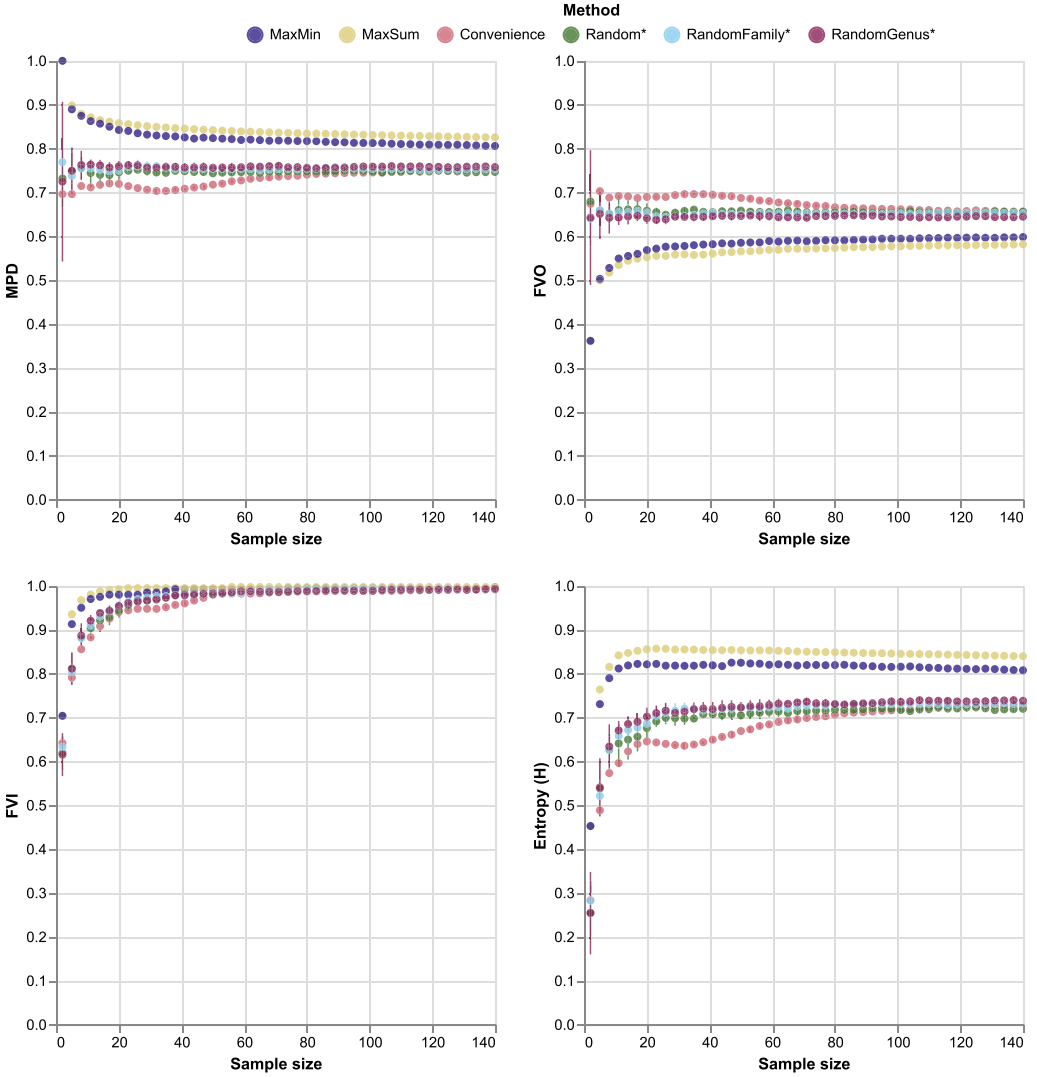


Figure 6 MPD (\uparrow), FVO (\downarrow), FVI (\uparrow), and H (\uparrow) for different sample sizes. Non-deterministic methods are indicated with an asterisk and averaged over 10 runs; error bars represent their standard deviation.

7. How Does Typological Diversity Relate to Generalizability?

In this work, we argue that increasing the linguistic representativeness of a language sample leads to better generalizability of results in multilingual NLP. Currently, multilingual language technology is typically only evaluated on a handful of seemingly randomly selected languages. This lack of systematic language sampling makes it difficult to assess how multilingually generalizable such technology really is. In this section, we assess whether sampling typologically diverse languages actually leads to better estimates of general multilingual performance. We apply our sampling framework to

Table 1

Results for $k = 20$ with all languages in Grambank as the frame. An asterisk indicates methods that are non-deterministic; these are averaged and the standard deviation over 10 random runs is listed.

Sampling Method	MPD \uparrow	FVO \downarrow	FVI \uparrow	H \uparrow
Convenience*	0.72 ± 0.00	0.69 ± 0.00	0.94 ± 0.00	0.65 ± 0.00
Random*	0.75 ± 0.01	0.66 ± 0.02	0.94 ± 0.02	0.68 ± 0.02
RandomFamily*	0.75 ± 0.01	0.65 ± 0.01	0.95 ± 0.01	0.69 ± 0.02
RandomGenus*	0.76 ± 0.02	0.64 ± 0.01	0.95 ± 0.01	0.70 ± 0.02
MaxSum	0.86	0.55	0.99	0.86
MaxMin	0.84	0.57	0.98	0.82

Table 2

Average results (accuracy) per sampling method ($k = 20$) on Taxi1500 dataset. For the non-deterministic sampling method, we report the average over 10 runs.

	<i>Non-deterministic Methods</i>			<i>Deterministic Methods</i>			
	Rand.	Rand. Fam	Rand. Gen.	Conv.	MaxSum	MaxMin	All
mBERT	14.3 ± 1.66	12.7 ± 1.22	13.5 ± 1.31	39.7	12.7	13.7	13.7
XLM-R-B	13.0 ± 2.91	9.7 ± 1.91	11.0 ± 2.81	56.8	7.7	13.8	11.0
XLM-R-L	12.4 ± 3.34	8.6 ± 2.05	10.3 ± 3.14	59.7	5.9	13.6	10.1
Glott500-m	16.8 ± 3.64	10.8 ± 1.70	12.9 ± 3.01	54.6	9.6	13.9	14.4

downstream NLP tasks, demonstrating how sampling diverse test languages affects the generalizability of the results. Specifically, we look into natural language understanding (topic classification), natural language generation (machine translation), and subword segmentation. Note that, as illustrated in Figure 1, the objective of systematic language sampling is to draw a language sample from the frame that optimally reflects the larger population. Thus, we evaluate the extent to which the distribution of results in typologically diverse samples reflects the distribution over a large group of languages.

7.1 Natural Language Understanding

To assess the extent to which samples drawn with different methods reflect general multilingual performance, it is important to sample from a dataset that includes sufficient typological variation. The Taxi1500 dataset (Ma et al. 2023), which spans more than 1,500 languages,¹⁰ covers a large part of the “typological space” (see Appendix B). Importantly, its evaluation set is parallel across languages, which facilitates cross-lingual comparison. The benchmarking task of Taxi1500 is classifying the topic of bible verses, where the six possible topics are *recommendation*, *faith*, *description*, *sin*, *grace*, and *violence*. These topics were manually annotated on an English bible, and applied to other languages through annotation projection to create a large-scale parallel benchmark. We compare the results of widely popular classification models (mBERT, Devlin et al. 2019; XLM-R, Conneau et al. 2020) as well as the more recently introduced Glott500-m (Imani et al. 2023).

First, in Table 2, we compare the average scores over the samples, and their similarity to the average over all languages. Ideally, diverse language sampling should lead to

¹⁰ The intersection between Grambank and Taxi1500 contains 1,506 languages.

averages that are more comparable to the average over all languages. We observe that this is indeed the case with our sampling methods, and that especially results obtained with the convenience baseline provide overly optimistic pictures of general multilingual performance, as compared to the more realistic performance estimates obtained with our carefully sampled language selections. Importantly, we highlight differences between MaxSum and MaxMin: MaxSum gives more conservative performance estimates than MaxMin, since, similar to variety sampling, it samples the most extreme cases. This may be useful to NLP practitioners, especially in scenarios where a worst-case estimate is desired.

In Figure 7, we go beyond average results, and inspect the *distribution* of the results. We compare the distribution of the accuracy over all 1,506 languages with samples ($k = 20$) from these methods: MaxSum, MaxMin (ours), and convenience (reflecting common practice in NLP). We observe that the convenience baseline gives consistent overestimations of multilingual performance, while the results of our systematic sampling method consistently reflect the distributions of results for all languages in the dataset. This implies that performance estimates on true typologically diverse languages indeed paint a more generalizable picture of the performance of multilingual models.

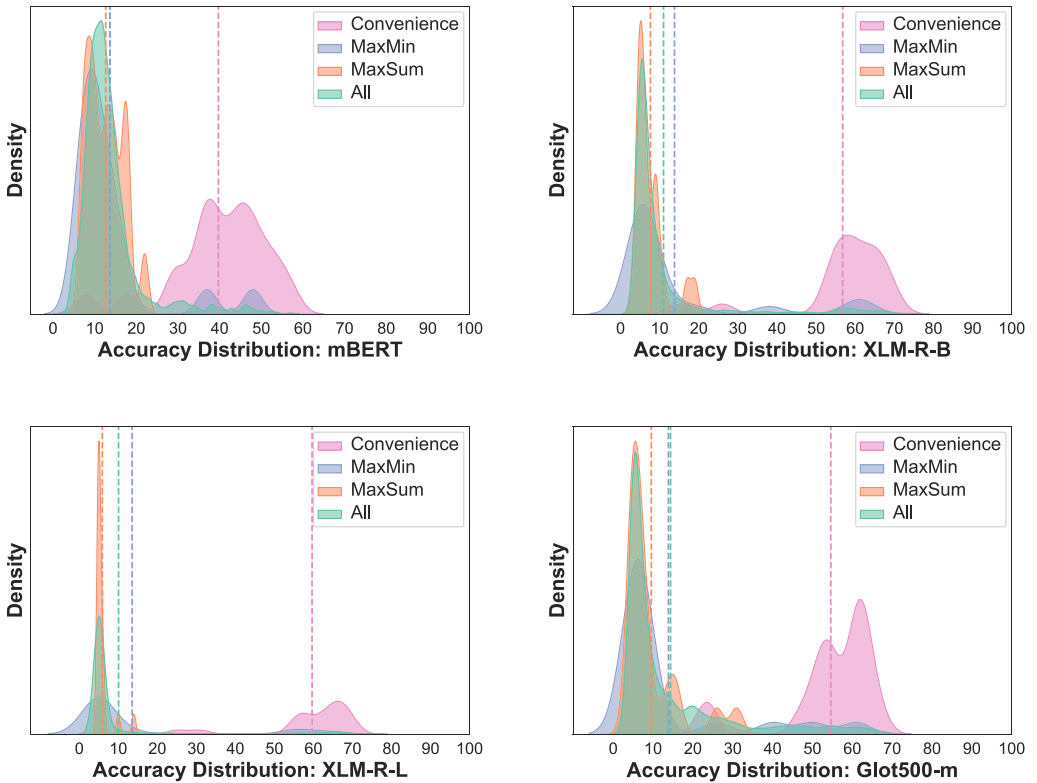


Figure 7 Accuracy of zero-shot performance distributions of language samples ($k = 20$) on the Taxi1500 evaluation set, contrasting all four deterministic sampling methods. Vertical lines represent averages per sampling method. Convenience sampling consistently over-estimates performance, as compared to MaxSum and MaxMin.

Table 3

Average results (BLEU) per sampling method ($k = 20$) on FloRes-200 dataset. Non-deterministic sampling methods are run 10 times.

<i>Non-deterministic Methods</i>			<i>Deterministic Methods</i>			
Rand.	Rand. Fam	Rand. Gen.	Conv.	MaxSum	MaxMin	All
34.74 \pm 3.29	32.74 \pm 1.18	33.00 \pm 0.98	46.64	30.05	31.52	32.69

Of course, inclusion or exclusion of languages in the training data for these models can be expected to affect the results. In Appendix C, we show the results when sampling only from the pre-training languages of each model. The same pattern holds: The convenience baseline systematically overestimates multilingual performance. Note that in our analysis, we investigated the zero-shot performance of the multilingual models, thus fine-tuning languages are not a confounding factor.

7.2 Natural Language Generation

Beyond natural language understanding, we evaluate the effect of diverse sampling on a generative task, namely, machine translation (MT). We evaluate the effect of the source language on MT performance to English, using the multilingual NLLB-200 model (3.3B parameters).¹¹ The results are retrieved from the OPUS Dashboard (Tiedemann and de Gibert 2023). Due to cross-lingual differences, MT performance can only be compared on the same target language and the same test sets. We choose to focus on investigating translation to English specifically, as this is the only language for which the dashboard contains results from all source languages in FloRes-200. English-centric MT has widely been prioritized in MT research, largely motivated by speaker numbers and information access, as argued by Bugliarello et al. (2020), among others. Furthermore, because we do not propose a general language “leaderboard”, but rather systematically investigate the influence of typological characteristics on NLP evaluation, we avoid the confounding factor of writing systems. We control for script in our experiments by sampling from languages that use the Latin script specifically, since this maximizes the number of languages in our sampling frame. It should be noted that this limits the generalizability of the findings of this experiment, since we exclude a sizable group of languages.

The results, as shown in Table 3 and Figure 8, demonstrate that convenience sampling gives an overestimation of performance of many-to-English MT. While performance may also be determined by the similarity to English, we see that MaxMin and MaxSum retrieve averages and distributions that resemble the distribution over the entire sampling frame. Furthermore, sampling with the MaxSum objective leads to a more conservative performance estimate than the MaxMin objective. Thus, MaxSum sampling is more suitable for application in worst-case evaluation scenarios, whereas the average of MaxMin samples can be expected to better reflect a general distribution.

7.3 Subword Segmentation

Subword tokenization (Sennrich, Haddow, and Birch 2016) is an important component in multilingual text processing, which is at the basis of many popular LLMs. While

¹¹ [huggingface/facebook/nllb-200-3.3B](https://huggingface.co/facebook/nllb-200-3.3B).




Figure 8 Distribution of results (BLEU) of many-to-English MT with NLLB-3.3B, on the languages in Latin script from the FloRes-200 dataset. Vertical lines represent averages.

it is not a downstream task, suboptimal tokenization can have large implications for downstream results. Splitting tokens into subwords facilitates better generalizability to languages with complex morphology, and better handles out-of-vocabulary tokens. Yet, tokenizers of popular multilingual models have been shown to retrieve varying results across languages with varying morphosyntactic properties (Gutierrez-Vasques et al. 2021; Gutierrez-Vasques, Bentz, and Samardžić 2023). Highly synthetic languages such as Finnish pose different tokenization challenges than languages with relatively little morphological complexity such as Dutch. The number of subwords can indicate over-segmentation, which can have far-reaching consequences for a user: for instance, the token-based cost of LLM APIs is higher, and processing more separate tokens introduces latency (Petrov et al. 2023). Previous work (Ahia et al. 2023; Petrov et al. 2023) therefore links segmentation to the concept of multilingual fairness. To extend the scope of our evaluation beyond strict downstream results to fairness more generally, we include a systematic comparison of subword segmenters.

Here, we conduct a large-scale cross-lingual sampling comparison of subword tokenization with tokenizers of popular multilingual models. For comparability between languages, the data should ideally be parallel across languages. Previous work (Ahia et al. 2023; Petrov et al. 2023) evaluates on the FLoRes-200 dataset, as this is multi-parallel. We considerably extend upon their language coverage in our analysis, by evaluating on text from the Parallel Bible Corpus (Mayer and Cysouw 2014), which is the largest massively multi-parallel dataset in terms of language coverage. We first match language ISO codes to Glottocodes. We then select all languages that have Grambank coverage, and control for script (i.e., we filter out non-Latin scripts). This retrieves 571 languages, which is the broadest language coverage thus far in tokenization analysis. This is important, because we aim to test to what extent the sample represents the sampling frame, and potentially the sampling universe. We select the longest bible per language, and sample 1,000 commonly available verses for our evaluation, to retain multi-parallelism. We analyze tokenizers of four popular language models with

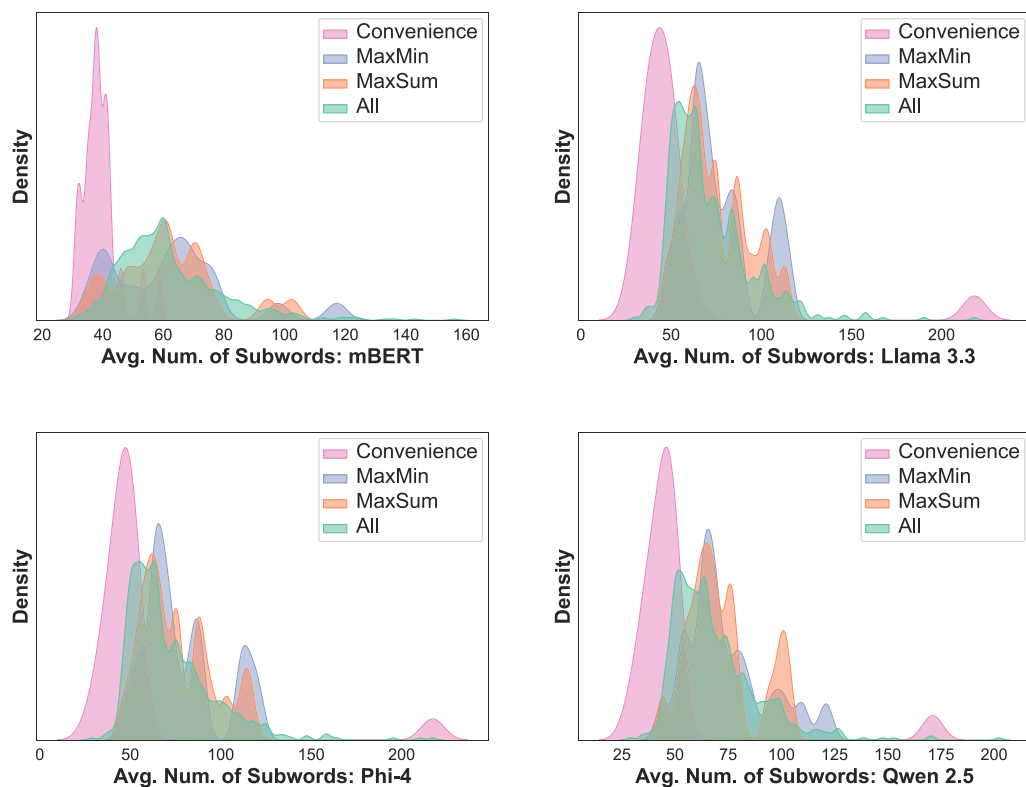


Figure 9 Average number of subwords per verse across four popular tokenizers, with different sampling strategies.

tokenizers publicly available on HuggingFace:¹² multilingual BERT (Devlin et al. 2019), Llama 3.3 70B Instruct (Dubey et al. 2024), Qwen2.5-32B (Team 2024), and Phi 4 (Abdin et al. 2024). Similar to Ahia et al. (2023), we measure the amount of segmentation through the average number of subwords per verse.¹³

We compare our sampling methods (MaxSum, MaxMin) with the only baseline that is deterministic for $k = 20$: convenience. From the 571 total languages, we sample 20 with each method. For each tokenizer, we compare the spread in average number of subwords per verse for all three sampling methods, and compare those with the spread for all 570 languages in our experiment. This gives us an estimate of how generalizable the sample of 20 languages is with respect to the total 500+ languages included in the experiment.

The results are shown in Figure 9. For all four models, we observe that the convenience baseline retrieves a considerably lower average number of subwords, with

¹² <https://huggingface.co>.

¹³ We do not use the fertility measure (Rust et al. 2021), because it assumes similar word tokenizer performance across languages. We do not use “premiums”, the disparity of tokenization length between parallel sentences in two languages (Petrov et al. 2023), because these are by definition relative to another language; we instead want to compare *across* languages directly.

a generally smaller range than other methods. This implies that evaluating on the 20 most commonly included languages in typologically diverse samples in NLP gives an overly optimistic image of *general* multilingual tokenizer performance. Yet, MaxMin retrieves averages and spreads that more closely resemble the average and spread of *all* 500+ languages in the experiment. This suggests that a priori typologically informed language sampling can improve the generalizability of the results, based on a language sample.

8. Further Use Cases

Our framework can be used in a variety of ways. In this section, we provide examples of practical use cases, to inspire future research with systematic language selection. Contrary to the experiments in Section 6, we now also deal with a corpus frame. This means that we do not merely sample from a typological database, but investigate more realistic data availability situations. Our experiments, frames, and samples are publicly available in our code repository.¹⁴

8.1 Guiding Dataset Expansion Efforts

Data availability is an obstacle for truly diverse sampling, as corpus frames may be limited for certain NLP tasks. At the same time, data collection and annotation efforts can be laborious and expensive. Thus, it may be useful to know beforehand how data collection for one language may impact the generalizability of an evaluation set. Our framework can be used to assess the diversity of existing benchmarks, inform future data collection efforts, and quantify the relative improvement in language diversity.

We present a use case with five popular existing multilingual evaluation benchmarks. They represent a range of tasks, including machine translation (FloRes-200; Team NLLB et al. 2022), dependency parsing (Universal Dependencies v2.14; Zeman et al. 2024), question answering (TyDiQA; Clark et al. 2020), commonsense reasoning (XCOPA; Ponti et al. 2020), and generative language modeling (Aya Evaluation Suite [human-annotated]; Singh et al. 2024). We focus on human-curated datasets specifically, as human annotations are typically more costly to gather than automatically generated data, and thus informed expansion is more relevant. The benchmarks vary in size and in the extent to which the included languages were carefully selected. For example, for FloRes-200 and Universal Dependencies (UD), there seem to be no explicit selection criteria, as expanding language coverage was the main objective. Yet, some other datasets were explicitly created with typological diversity in mind. TyDiQA was created with the aim to include typologically diverse languages. However, in line with our findings in Section 2, the authors do not specify systematic sampling criteria and only post-hoc mention typological features to “highlight the breadth of phenomena” of the included languages. The authors of XCOPA explicitly aim for variety sampling, but only seem to measure the typological diversity of their sample post-hoc.

Here, we quantify the diversity of each benchmark dataset, and use our framework to assess which expansion language would most increase typological diversity, and how this impacts the total set. To this end, we provide a starting sample, which is the intersection of the languages in the datasets, and those in Grambank. For retrieving

¹⁴ https://github.com/esther2000/typdiv-sampling/use_cases.

Table 4

The expansion language that most increases typological diversity for five popular multilingual benchmarks, as retrieved with MaxSum sampling.

Dataset	$ L (L \cap GB)$	H	FVI	+ Language	H'	FVI'
FloRes-200 (Team NLLB et al. 2022)	195 (105)	0.717	0.988	Tariana	0.721	0.988
UD v2.14 (Zeman et al. 2024)	158 (80)	0.681	0.985	Tariana	0.687	0.985
TyDiQA (Clark et al. 2020)	11 (7)	0.627	0.883	Movima	0.693	0.928
XCOPA (Ponti et al. 2020)	11 (7)	0.599	0.873	Tariana	0.667	0.913
Aya Evaluation Suite (human-annotated) (Singh et al. 2024)	7 (5)	0.571	0.841	Yele	0.660	0.898

the next-best language, we then choose a sample size of the starting sample, plus one. Note that this number can be raised in future applications. We sample with the MaxSum objective, as this is effective in increasing the total diversity (Section 6.3). We report diversity based on entropy (H) and feature value inclusion (FVI), as discussed in Section 6.2.

The results (Table 4) show that adding the next-best languages to the existing datasets always increases total diversity in terms of entropy. This increase is relatively small for the larger datasets (FloRes-200, UD), but larger for the smaller datasets (TyDiQA, XCOPA, Aya Evaluation Suite), where the number of added languages is a larger proportion of the total. FVI is not increased by adding a single language to the larger datasets. This is unsurprising, as the value was already nearly maximal. This suggests that including as many languages as possible is a good strategy for maximizing the included features, as argued for in variety sampling (Miestamo, Bakker, and Arppe 2016).

Interestingly, one language stands out as the most diverse expansion language for multiple datasets (FloRes-200, UD, and XCOPA): Tariana (language family: Arawakan). From a linguistic typology perspective, this makes sense. Firstly, Tariana is spoken in “a very remote area” in the Vaupés area in the northwest of Brazil, which is “not easy to get to” (Aikhenvald 2003b). Also, Tariana is a polysynthetic language (Aikhenvald 2003a), a grammatical property that is likely uncommon in popular NLP datasets. Yet, there are more extensive, cultural reasons for Tariana standing out. The area where the language is spoken is characterized by “obligatory multilingualism, dictated by the principles of linguistic exogamy” (Aikhenvald 2003a). This means that marriage only occurs between speakers of different languages. There has been a “strong inhibition against ‘language-mixing’, viewed in terms of lexical loans”, and the language includes “independent innovations [...] divergent from those found in closely related languages” (Aikhenvald 2003b). While this case is an indicator that our method indeed achieves high diversity and thus does what we expect, we do not argue that Tariana should then be included in all NLP datasets. Instead, speaker needs and data availability should be taken into account, which can be done by narrowing the sampling frame. We demonstrate such a case in the next section.

Case Study: Universal Dependencies. So far we have used all (cropped) languages in Grambank as the sampling frame. In practice, this may not be realistic, as one might want to take into account speaker needs and annotator availability. The sampling frame can then be smaller. To analyze our framework in such a scenario, we narrow the

Table 5

Effects in diversity metrics from adding Seri to UD v2.14.

MPD \uparrow	MPD' \uparrow	FVO \downarrow	FVO' \downarrow	FVI \uparrow	FVI' \uparrow	H \uparrow	H' \uparrow
0.725	<u>0.728</u>	0.679	<u>0.677</u>	<u>0.985</u>	<u>0.985</u>	0.681	<u>0.685</u>

sampling frame for UD to the languages in their *Possible Future Extensions* list.¹⁵ These are “[languages for which] people have expressed interest in providing annotated data [...], but [for which] no valid data has been provided so far”. We manually annotate the Glottocodes for these languages, and find that the intersection of previously not included Glottocodes and Grambank contains 17 languages. These languages are the sampling frame for finding the “next best” language. We apply MaxSum sampling, which retrieves as next best extension language: Seri. Again, this makes sense from a linguistic typology perspective, as Seri is a language isolate which “does indeed have some very special characteristics” (Marlett 2000). In Table 5, we report the effect of adding Seri to UD v2.14 on all our metrics. We observe that, with the exception of FVI, diversity is improved with respect to all metrics.

8.2 Other Distance Maximization

Beyond typological features, our method serves primarily as a *framework*, which can be extended to other features that describe languages. For example, it can be used to sample languages that are spoken in geographically distant areas. Such language sampling has been an important methodological focus in linguistic typology. Similar to sampling from genealogical groupings, geographically diverse sampling is commonly performed with pre-defined groupings, such as macroareas (Dryer 1989, 1992; Hammarström and Donohue 2014) or the finer-grained AUTOTYP areas (Nichols and Bickel 2009). In Section 3, we established that the usefulness of phylogenetic language groupings can be hindered by their lack of granularity. The same applies to geographical language groupings. Firstly, all languages from one such area (e.g., “Eurasia”) are seen as being equally close. Secondly, languages that are geographically close, but not in the same macroarea, are not seen as similar. Most previous work in typology that uses absolute geographical distance instead of groupings (Jaeger et al. 2011; Cysouw 2013; Bjerva et al. 2020) manually defines a set threshold, for instance of 1,000 kilometers.¹⁶

Our framework allows for diverse geographical sampling, directly from distances, without the need to manually define a distance threshold. We further demonstrate this use case here. We first retrieve language coordinates from Grambank, which corresponds to step 1 in our framework. We use these to calculate the absolute distance in kilometers between any two coordinates (step 2). These pairwise distances are then used for sampling geographically diverse languages. Specifically, we use the MaxMin objective, with the objective to maximize the minimum distance between any two points. Figure 10 shows the resulting sample when selecting twenty languages.

¹⁵ As listed on the homepage at the time of writing: <https://universaldependencies.org>.

¹⁶ Still, it should be noted that distance-based stratification does not take into account long-distance contact, and that reducing languages to coordinates is not necessarily accurate.

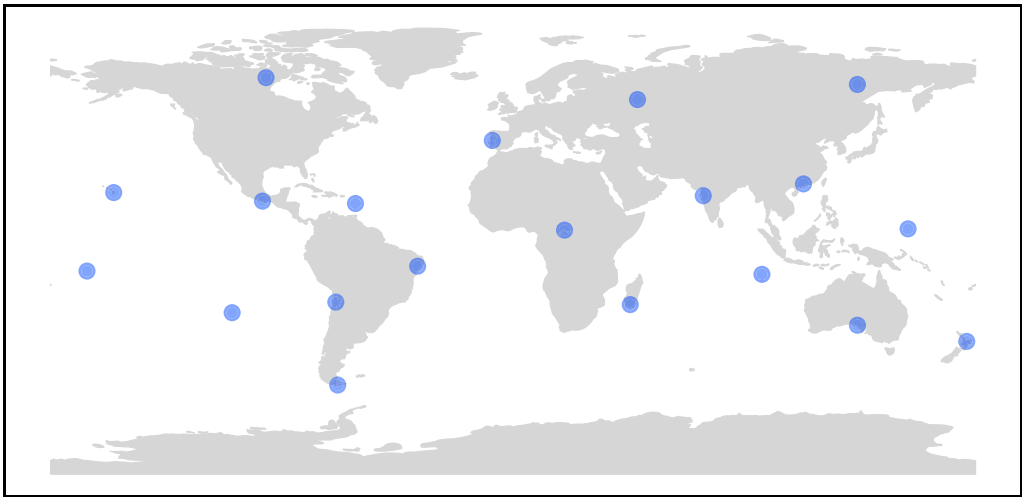


Figure 10

We use MaxMin sampling to find twenty geographically distant languages, based on coordinate pairs.

9. Limitations, Ethical Considerations, and Discussion

Our framework enables typologically diverse language sampling. Critical assessment of language diversity is vital for the evaluation and development of multilingual language technology that is fair across languages. Still, it should be noted that generalizability across typological characteristics constitutes only a fragment of multilingually fair NLP. For instance, Grambank only addresses grammatical phenomena. Such typological databases do not provide cultural information, which may be key for certain research questions (Hershcovich et al. 2022).

Moreover, the feature coverage in typological databases such as Grambank is incomplete, which should be taken into account when drawing conclusions. As such, our evaluation is not completely representative. However, since bibliographic bias in NLP research tends to be much stronger than in typological databases (see convenience baseline, Section 6.3), we believe that information from typological databases can actually enable informed expansion of language coverage. For example, through informed dataset expansion (Section 8.1), we hope our framework can be used to further increase typological representativeness in future datasets.

Lastly, we acknowledge that reducing languages to coordinates or points in multidimensional space is by default simplistic. Languages are more than objects of study: They are central to human communication and inherently involve humans. As such, we urge researchers to consider factors beyond typological diversity in language sampling or dataset expansion. Instead of sampling *only* based on typological diversity, we emphasize the importance of incorporating a human-centered perspective. While developing more generalizable multilingual NLP tools has the potential of mitigating unequal access to language technology, this should be a collaborative effort that involves speakers (Bird 2020).

10. Conclusions

In this work, we systematically analyze common language sampling strategies in NLP, and find that these are insufficient for typologically diverse language sampling. Guided by research in linguistic typology, we propose two sampling algorithms. We compare the samples obtained by our methods with strong baselines. These samples are evaluated with four typological diversity metrics that show that our method consistently retrieves language samples with higher typological diversity. For downstream tasks, we show that our sampling algorithms provide a much more realistic estimate of generalizability as compared to, e.g., convenience sampling, thereby highlighting the importance of addressing typological diversity in the context of empirical NLP work. While we focus on achieving high typological distance, our framework can be used with *any* type of information that one wishes to make “diverse” generalizable claims about. Furthermore, our method can also be used for finding typologically *similar* languages, e.g., with use cases in transfer learning scenarios. We recommend that future work in NLP aiming for typological generalizability use our informed selection method (while not losing sight of the human-centered perspective).

Appendix A. Sampling Evaluation by k for Intersection of Sampling Frames

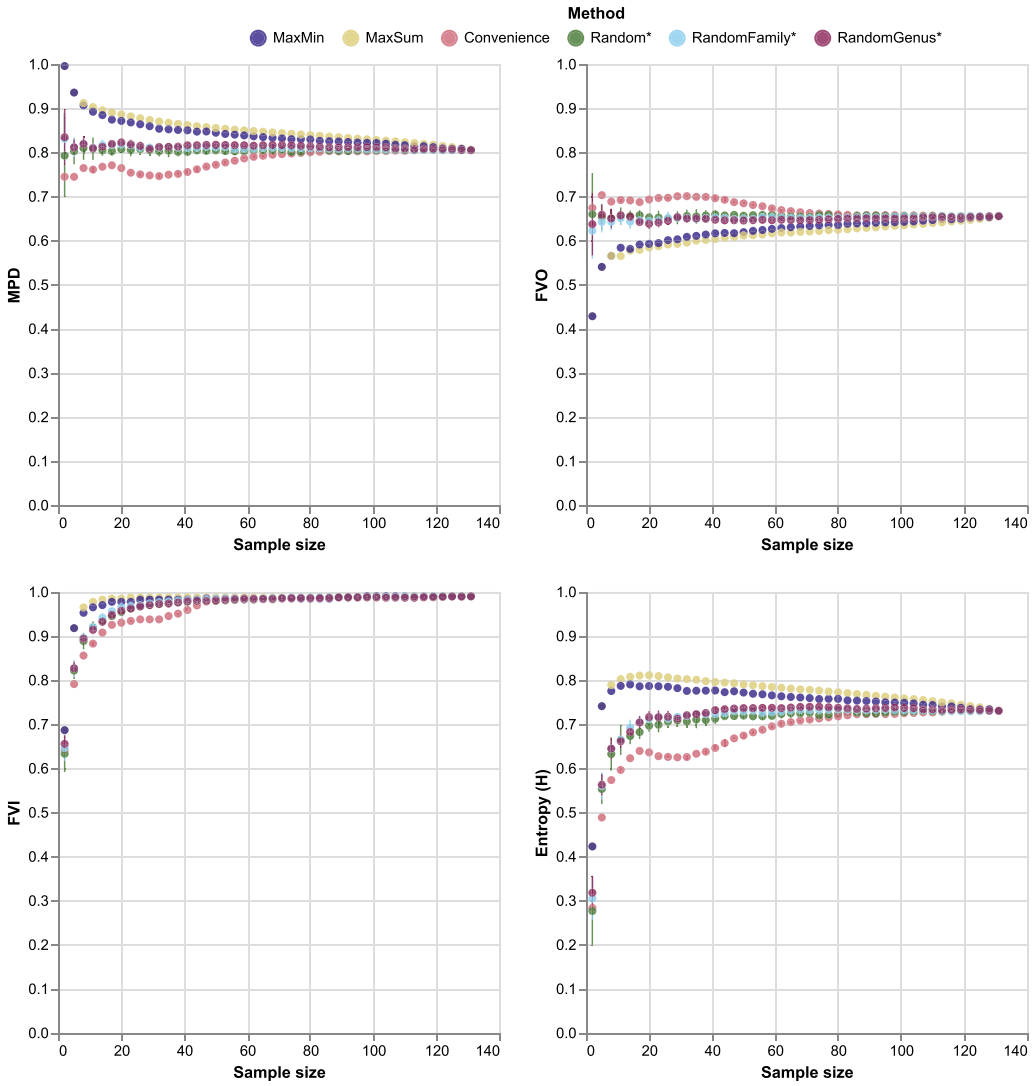


Figure A.1
 MPD (\uparrow), FVO (\downarrow), FVI (\uparrow), and H (\uparrow) for different sample sizes where the sampling frame is equal (intersection for all methods). Non-deterministic methods are indicated with an asterisk and averaged over 10 runs, error bars represent their standard deviation. Colors are from Paul Tol's color-blind safe muted qualitative scheme.

Appendix B. Diversity of Downstream Sampling Frames

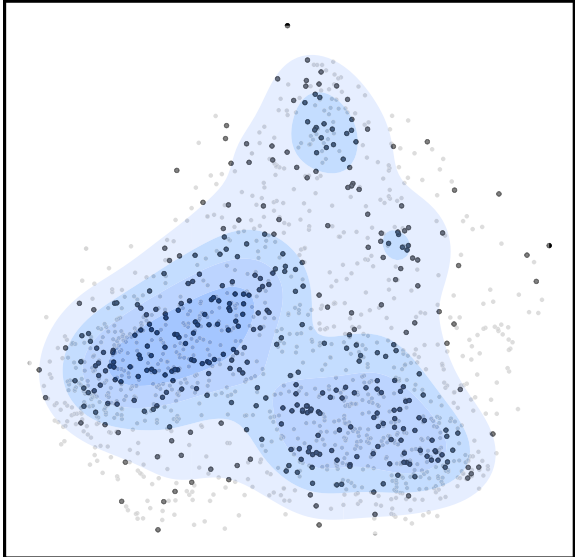


Figure B.1
PCA plot illustrating the extent to which Taxi1500 covers Grambank’s typological space.




Figure B.2
PCA plot illustrating the extent to which FloRes-200 covers Grambank’s typological space.

Appendix C. Downstream Results on Pre-training Languages Only

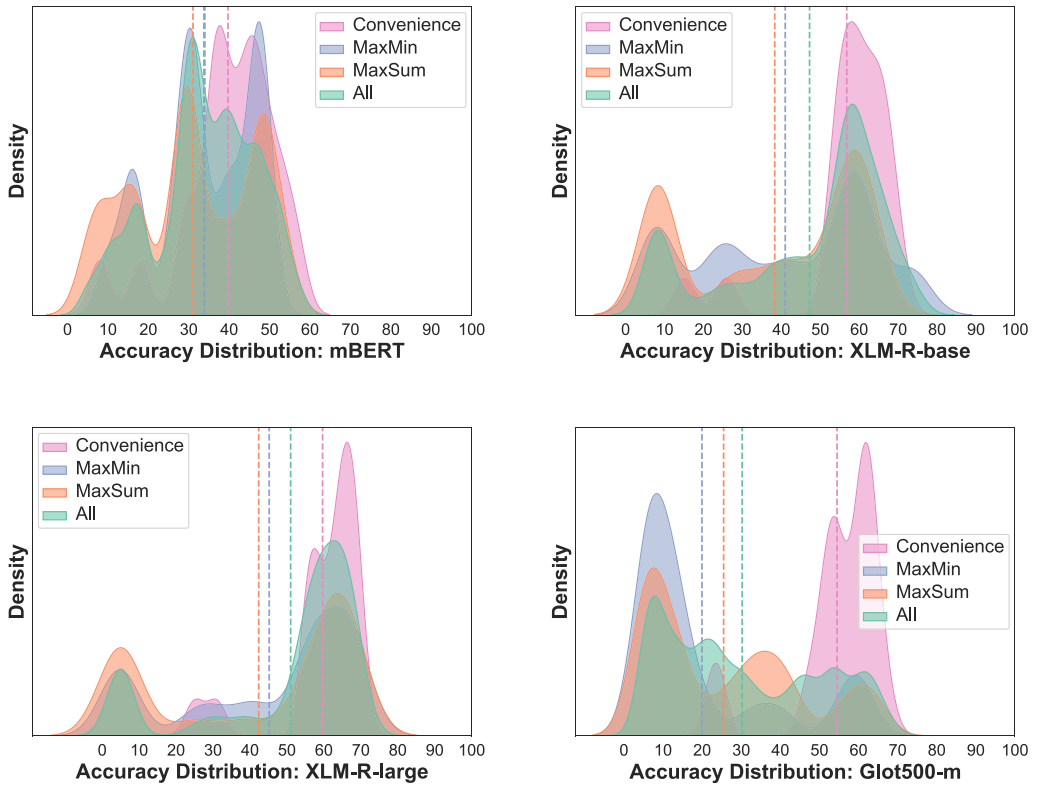


Figure C.1 Accuracy distributions zero-shot performance of language samples ($k = 20$) on the Taxi1500 evaluation set, contrasting all four deterministic sampling methods, when sampling only from the pre-training languages of each model.

Acknowledgments

We thank the members of LAGoM-NLP at KU Leuven, the AAU-NLP lab at Aalborg University, and the Helsinki-NLP group at the University of Helsinki for feedback on earlier versions of this paper. We thank Robert Östling for pointing us to Anna Sjöberg's thesis. We thank Kaius Sinnemäki and Anna Sjöberg for discussing the scope of this project with us from a typological perspective. We thank Hedvig Skirgård for helpful feedback on cropping the dataset according to coverage, removing macrolanguages, and the overall scope of this paper. Any remaining errors are our own. E.P. and J.B. are funded by the Carlsberg Foundation, under the *Semper Ardens: Accelerate* programme (project nr. CF21-0454). W.P. is funded by a KU Leuven Bijzonder Onderzoeksfonds C1 project with reference C14/23/096. This work received support from the CA21167 COST action UniDive, funded by COST (European Cooperation in Science and Technology).

References

- Abdin, Marah, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, et al. 2024. Phi-4 technical report. <https://arxiv.org/abs/2412.08905>
- Ács, Judit, Ákos Kádár, and Andras Kornai. 2021. Subword pooling makes a difference. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2284–2295. <https://doi.org/10.18653/v1/2021.eacl-main.194>
- Ahia, Orevaoghene, Sachin Kumar, Hila Gonen, Jungo Kasai, David Mortensen, Noah Smith, and Yulia Tsvetkov. 2023. Do all languages cost the same? Tokenization in the era of commercial language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9904–9923. <https://doi.org/10.18653/v1/2023.emnlp-main.614>
- Aikhenvald, Alexandra Y. 2003a. The language and its speakers. In *A Grammar of Tariana, from Northwest Amazonia*, Cambridge Grammatical Descriptions. Cambridge University Press, pages 1–24. <https://doi.org/10.1017/CB09781107050952.005>
- Aikhenvald, Alexandra Y. 2003b. Teaching Tariana, an endangered language from northwest Amazonia. *International Journal of the Sociology of Language*, 2003(161):125–139. <https://doi.org/10.1515/ijsl.2003.025>
- Bakker, Dik. 2010. Language sampling. *Oxford handbooks in linguistics*. <https://doi.org/10.1093/oxfordhb/9780199281251.013.0007>
- Bell, Alan. 1978. Language samples. *Universals of Human Language*, 1:123–156.
- Bender, Emily M. 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. <https://doi.org/10.3115/1642038.1642044>
- Bender, Emily M. 2011. On achieving and evaluating language-independence in NLP. *Linguistic Issues in Language Technology*, 6. <https://doi.org/10.33011/lilt.v6i.1239>
- Bird, Steven. 2020. Decolonising speech and language technology. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3504–3519. <https://doi.org/10.18653/v1/2020.coling-main.313>
- Bjerva, Johannes, Elizabeth Salesky, Sabrina J. Mielke, Aditi Chaudhary, Giuseppe G. A. Celano, Edoardo Maria Ponti, Ekaterina Vylomova, Ryan Cotterell, and Isabelle Augenstein. 2020. SIGTYP 2020 shared task: Prediction of typological features. In *Proceedings of the Second Workshop on Computational Research in Linguistic Typology*, pages 1–11. <https://doi.org/10.18653/v1/2020.sigtyp-1.1>
- Bugliarello, Emanuele, Sabrina J. Mielke, Antonios Anastasopoulos, Ryan Cotterell, and Naoaki Okazaki. 2020. It's easier to translate out of English than into it: Measuring neural translation difficulty by cross-mutual information. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1640–1649. <https://doi.org/10.18653/v1/2020.acl-main.149>
- Clark, Jonathan H., 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. https://doi.org/10.1162/tac1_a_00317

- Conneau, Alexis, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451. <https://doi.org/10.18653/v1/2020.acl-main.747>
- Cysouw, Michael. 2013. Disentangling geography from genealogy. In *Space in Language and Linguistics: Geographical, Interactional, and Cognitive Perspectives*. de Gruyter. <https://doi.org/10.1515/9783110312027.21>
- Dahl, Östen. 2008. An exercise in a posteriori language sampling. *Language Typology and Universals*, 61(3):208–220. <https://doi.org/10.1524/stuf.2008.0021>
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- Dixon, John K. 1979. Pattern recognition with partly missing data. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(10):617–621. <https://doi.org/10.1109/TSMC.1979.4310090>
- Dixon, Robert M. W. 1997. *The Rise and Fall of Languages*. Cambridge University Press. <https://doi.org/10.1017/CB09780511612060>
- Dryer, Matthew S. 1989. Large linguistic areas and language sampling. *Studies in Language. International Journal sponsored by the Foundation "Foundations of Language"*, 13(2):257–292. <https://doi.org/10.1075/s1.13.2.03dry>
- Dryer, Matthew S. 1992. The Greenbergian word order correlations. *Language*, 68(1):81–138. <https://doi.org/10.1353/lan.1992.0028>
- Dryer, Matthew S. 1998. Why statistical universals are better than absolute universals. In *Papers from the 33rd Regional Meeting of the Chicago Linguistic Society*, pages 1–23.
- Dryer, Matthew S. and Martin Haspelmath. 2013. *WALS Online*. Leipzig: Max Planck Institute for Evolutionary Anthropology.
- Dubey, Abhimanyu, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Georgi, Ryan, Fei Xia, and William Lewis. 2010. Comparing language similarity across genetic and typologically-based groupings. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 385–393.
- Greenberg, Joseph Harold. 1963. Some universals of grammar with particular reference to the order of meaningful elements. In *Universals of Language*. MIT Press, pages 40–70.
- Gutierrez-Vasques, Ximena, Christian Bentz, and Tanja Samardžić. 2023. Languages through the looking glass of BPE compression. *Computational Linguistics*, 49(4):943–1001. <https://doi.org/10.1162/coli.a.00489>
- Gutierrez-Vasques, Ximena, 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. <https://doi.org/10.18653/v1/2021.eacl-main.302>
- Guzmán Naranjo, Matías and Laura Becker. 2022. Statistical bias control in typology. *Linguistic Typology*, 26(3):605–670. <https://doi.org/10.1515/lingty-2021-0002>
- Hammarström, Harald and Mark Donohue. 2014. Some principles on the use of macro-areas in typological comparison. *Language Dynamics and Change*, 4(1):167–187. <https://doi.org/10.1163/22105832-00401001>
- Hammarström, Harald, Robert Forkel, Martin Haspelmath, and Sebastian Bank. 2024. *glottolog/glottolog: Glottolog database 5.0*. <https://doi.org/10.5281/zenodo.10804357>
- Haynie, Hannah J., Damián Blasi, Hedvig Skirgård, Simon J. Greenhill, Quentin D. Atkinson, and Russell D. Gray. 2023. Grambank’s typological advances support computational research on diverse languages. In *Proceedings of the 5th Workshop on Research in Computational Linguistic Typology and Multilingual NLP*, pages 147–149. <https://doi.org/10.18653/v1/2023.sigtyp-1.17>

- Hershcovich, Daniel, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza Fierro, Katerina Margatina, Phillip Rust, and Anders Søgaard. 2022. Challenges and strategies in cross-cultural NLP. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6997–7013. <https://doi.org/10.18653/v1/2022.acl-long.482>
- Imani, Ayyoob, Peiqin Lin, Amir Hossein Kargaran, Silvia Severini, Masoud Jalili Sabet, Nora Kassner, Chunlan Ma, Helmut Schmid, André Martins, François Yvon, and Hinrich Schütze. 2023. Glot500: Scaling multilingual corpora and language models to 500 languages. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1082–1117. <https://doi.org/10.18653/v1/2023.acl-long.61>
- Jaeger, T. Florian, Peter Graff, William Croft, and Daniel Pontillo. 2011. Mixed effect models for genetic and areal dependencies in linguistic typology. *Linguistic Typology*, 15(2):281–319. <https://doi.org/10.1515/lity.2011.021>
- Kashyap, Abhishek Kumar. 2019. Language typology. *The Cambridge Handbook of Systemic Functional Linguistics*. pages 767–792. <https://doi.org/10.1017/9781316337936.031>
- Kaufman, Leonard and Peter J. Rousseeuw. 2009. *Finding Groups in Data: An Introduction to Cluster Analysis*. John Wiley & Sons.
- Kuo, Ching Chung, Fred Glover, and Krishna S. Dhir. 1993. Analyzing and modeling the maximum diversity problem by zero-one programming. *Decision Sciences*, 24(6):1171–1185. <https://doi.org/10.1111/j.1540-5915.1993.tb00509.x>
- Littell, Patrick, 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. <https://doi.org/10.18653/v1/E17-2002>
- Ma, Chunlan, Ayyoob ImaniGooghari, Haotian Ye, Ehsaneddin Asgari, and Hinrich Schütze. 2023. Taxi1500: A multilingual dataset for text classification in 1500 languages. *arXiv preprint arXiv:2305.08487*. <https://doi.org/10.21203/rs.3.rs-3235946/v1>
- Majewska, Olga, 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. <https://doi.org/10.18653/v1/2020.coling-main.423>
- Marlett, Stephen A. 2000. Why the Seri language is important and interesting. *Journal of the Southwest*, pages 611–633.
- Martí, Rafael, Micael Gallego, Abraham Duarte, and Eduardo G. Pardo. 2013. Heuristics and metaheuristics for the maximum diversity problem. *Journal of Heuristics*, 19:591–615. <https://doi.org/10.1007/s10732-011-9172-4>
- Mayer, Thomas and Michael Cysouw. 2014. Creating a massively parallel Bible corpus. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 3158–3163.
- Miestamo, Matti, Dik Bakker, and Antti Arppe. 2016. Sampling for variety. *Linguistic Typology*, 20(2):233–296. <https://doi.org/10.1515/lingty-2016-0006>
- Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems*, 26.
- Nichols, Johanna and Balthasar Bickel. 2009. The AUTOTYP genealogy and geography database: 2009 release. URL: <https://github.com/autotyp/autotyp-data>
- Parreño, Francisco, Ramón Álvarez-Valdés, and Rafael Martí. 2021. Measuring diversity. A review and an empirical analysis. *European Journal of Operational Research*, 289(2):515–532. <https://doi.org/10.1016/j.ejor.2020.07.053>
- Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Perkins, Revere D. 1988. The covariation of culture and grammar. *Studies in Syntactic*

- Typology*, pages 359–378. <https://doi.org/10.1075/ts1.17.22per>
- Petrov, Aleksandar, Emanuele La Malfa, Philip Torr, and Adel Bibi. 2023. Language model tokenizers introduce unfairness between languages. In *Advances in Neural Information Processing Systems*, 36:36963–36990.
- Pikuliak, Matúš and Marian Simko. 2022. Average is not enough: Caveats of multilingual evaluation. In *Proceedings of the 2nd Workshop on Multi-lingual Representation Learning (MRL)*, pages 125–133. <https://doi.org/10.18653/v1/2022.mrl-1.13>
- Ploeger, Esther, Wessel Poelman, Miryam de Lhoneux, and Johannes Bjerva. 2024. What is “typological diversity” in NLP? In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5681–5700. <https://doi.org/10.18653/v1/2024.emnlp-main.326>
- Ponti, Edoardo Maria, 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. <https://doi.org/10.18653/v1/2020.emnlp-main.185>
- Rijkhoff, Jan and Dik Bakker. 1998. Language sampling. *Linguistic Typology*, 2(3):263–314. <https://doi.org/10.1515/lity.1998.2.3.263>
- Rijkhoff, Jan, 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. <https://doi.org/10.1075/sl.17.1.07rij>
- Rust, Phillip, 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. <https://doi.org/10.18653/v1/2021.ac1-long.243>
- Sennrich, Rico, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725. <https://doi.org/10.18653/v1/P16-1162>
- Singh, Shivalika, Freddie Vargus, Daniel Dsouza, Börje F. Karlsson, Abinaya Mahendiran, Wei-Yin Ko, Herumb Shandilya, Jay Patel, Deividas Mataciunas, Laura O'Mahony, and others. 2024. Aya Dataset: An open-access collection for multilingual instruction tuning. *arXiv preprint arXiv:2402.06619*. <https://doi.org/10.18653/v1/2024.ac1-long.620>
- Sjöberg, Anna. 2023. *Knowledge Predication: A Semantic Typology*. Ph.D. thesis, Department of Linguistics, Stockholm University.
- Skirgård, Hedvig, Hannah J. Haynie, Damián E. Blasi, Harald Hammarström, Jeremy Collins, Jay J. Latache, Jakob Lesage, Tobias Weber, Alena Witzlack-Makarevich, Sam Passmore, et al. 2023. Grambank reveals the importance of genealogical constraints on linguistic diversity and highlights the impact of language loss. *Science Advances*, 9(16):eadg6175. <https://doi.org/10.1126/sciadv.adg6175>, PubMed: 37075104
- Skirgård, Hedvig, Hannah J. Haynie, Harald Hammarström, Damián E. Blasi, Jeremy Collins, Jay Latache, Jakob Lesage, Tobias Weber, Alena Witzlack-Makarevich, Michael Dunn, et al. 2023. Grambank v1.0. Dataset. <https://doi.org/10.5281/zenodo.7740140>
- Stoll, Sabine and Balthasar Bickel. 2013. Capturing diversity in language acquisition research. *Language Typology and Historical Contingency*. pages 195–216. <https://doi.org/10.1075/ts1.104.08slo>
- Team, Qwen. 2024. Qwen2.5: A party of foundation models. <https://qwenlm.github.io/blog/qwen2.5/>
- Team NLLB, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, et al. 2022. No language left behind: Scaling human-centered machine translation. *arXiv preprint arXiv:2207.04672*. <http://arxiv.org/abs/2207.04672>
- Tiedemann, Jörg and Ona de Gibert. 2023. The OPUS-MT dashboard – A toolkit for a systematic evaluation of open machine translation models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 3: System*

- Demonstrations*), pages 315–327. <https://doi.org/10.18653/v1/2023.acl-demo.30>
- Touvron, Hugo, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. LLaMa: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Wichmann, Søren and Eric W. Holman. 2010. Pairwise comparisons of typological profiles. In Jan Wohlgemuth and Michael Cysouw, editors, *Rethinking Universals: How Rarities Affect Linguistic Theory*. De Gruyter Mouton, pages 241–254. <https://doi.org/10.1515/9783110220933.241>
- Xu, Hongzhi, 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. <https://doi.org/10.18653/v1/2020.acl-main.596>
- Yadavalli, Aditya, 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. <https://doi.org/10.18653/v1/2023.acl-long.657>
- Yuan, Fei, Shuai Yuan, Zhiyong Wu, and Lei Li. 2023. How multilingual is multilingual LLM? *arXiv preprint arXiv:2311.09071*.
- Zeman, Daniel, Joakim Nivre, Mitchell Abrams, Elia Ackermann, Noëmi Aeppli, Hamid Aghaei, Željko Agić, Amir Ahmadi, Lars Ahrenberg, Chika Kennedy Ajede, et al. 2024. Universal dependencies 2.14. LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Zhang, Xiang, Senyu Li, Bradley Hauer, Ning Shi, and Grzegorz Kondrak. 2023. Don't trust ChatGPT when your question is not in English: A study of multilingual abilities and types of LLMs. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7915–7927. <https://doi.org/10.18653/v1/2023.emnlp-main.491>