

Genre as Weak Supervision for Cross-lingual Dependency Parsing

Max Müller-Eberstein and Rob van der Goot and Barbara Plank

Department of Computer Science

IT University of Copenhagen, Denmark

mamy@itu.dk, robv@itu.dk, bapl@itu.dk

Abstract

Recent work has shown that monolingual masked language models learn to represent data-driven notions of language variation which can be used for domain-targeted training data selection. Dataset *genre* labels are already frequently available, yet remain largely unexplored in cross-lingual setups. We harness this genre metadata as a weak supervision signal for targeted data selection in zero-shot dependency parsing. Specifically, we project treebank-level *genre* information to the finer-grained sentence level, with the goal to amplify information implicitly stored in unsupervised contextualized representations. We demonstrate that genre is recoverable from multilingual contextual embeddings and that it provides an effective signal for training data selection in cross-lingual, zero-shot scenarios. For 12 low-resource language treebanks, six of which are test-only, our genre-specific methods significantly outperform competitive baselines as well as recent embedding-based methods for data selection. Moreover, genre-based data selection provides new state-of-the-art results for three of these target languages.

1 Introduction

Multilingual masked language models (MLMs) trained on immense quantities of heterogeneous texts (Devlin et al., 2019; Brown et al., 2020; Conneau et al., 2020) have recently made applications such as highly cross-lingual dependency parsing a reality (Kondratyuk and Straka, 2019). Adjacently, it has also been recognized that they capture characteristics relevant for training data selection (Aharoni and Goldberg, 2020) and can be efficiently fine-tuned for higher task-specific performance (Gururangan et al., 2020; Dai et al., 2020; Lauscher et al., 2020; Üstün et al., 2020). These considerations are especially important in computationally restricted environments and when data from the target distribution are unavailable.

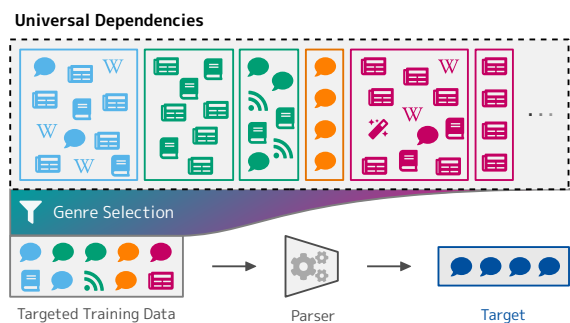


Figure 1: **Genre-driven Training Data Selection** for a zero-shot target treebank. In absence of annotated in-language data, we propose *genre* as a weak supervision signal for targeted instance selection from a large pool of out-of-language treebanks.

Universal Dependencies (Nivre et al., 2020; UD) provides an extensive testing ground for such scenarios: Its language diversity is constantly increasing (from 10 in v1.0 to 104 in v2.7) and low-resource languages are often limited to a single test-set-only treebank. As most of the 7,000+ languages in the world similarly lack any annotated training data, effective zero-shot transfer learning is crucial for achieving wider linguistic coverage.

Criteria for selecting training data within such settings vary, and a practitioner may determine relevance by proxy of language relatedness or treebank content. This leads us to the question: If our goal is to develop a parser for a known domain in an unseen language, can a signal such as *genre* guide our selection of cross-lingual training data from a significantly larger, diverse pool (Figure 1)?

Within the heterogeneity of written and spoken (transcribed) data, genre broadly encompasses variation along the functional role of a text (Kessler et al., 1997). A clear definition is complex if not impossible and communities refer to genre, domain, style or register in different ways (Kessler et al., 1997; Lee, 2001; Webber, 2009; Plank, 2011). In this work, we take a pragmatic approach and use

genre as defined by the 18 community-provided categories in UD (Zeman et al., 2020). These genres are assigned at the treebank level, and “are neither mutually exclusive nor based on homogeneous criteria, but [are] currently the best documentation that can be obtained” (Nivre et al., 2020).

Contributions In order to facilitate finer-grained, instance-level data selection for cross-lingual parsing in absence of in-language training data, we provide three contributions:

First, we provide an analysis of the genre distribution in UD v2.7 (Zeman et al., 2020) across 104 languages and 177 treebanks (Section 3).

Next, we introduce three targeted data selection strategies which amplify existing genre information in multilingual contextualized embeddings in order to enable sentence-level selection based on UD’s treebank-level genre annotations (Section 4).

Finally, we apply the extracted genre information to proxy training data selection for 12 typologically diverse low-resource treebanks. In absence of any in-language training data, our approach outperforms selection using treebank metadata alone as well as purely embedding-based instance selection and surpasses state-of-the-art results on three treebanks (Section 5).¹

2 Related Work

Despite advances in zero-shot performance (Devlin et al., 2019; Brown et al., 2020) and increasingly cross-lingual parsers (Kondratyuk and Straka, 2019), fine-tuning has remained a crucial step for achieving state-of-the-art performance. Meechan-Maddon and Nivre (2019) demonstrate that this holds true for low-resource languages in particular, with 200 training instances in the target or related languages producing better results on dependency parsing than a model trained on all available data. Lauscher et al. (2020) further show that as few as 10 samples in the target language can double parsing performance. Üstün et al. (2020) propose UDapters, which integrate language and task-specific adaptation modules into the parser to improve cross-lingual, zero-shot performance.

Considering factors complementary to language is equally important: MLMs can for instance be improved for specific domains such as Twitter or medical texts by fine-tuning on the same or related sources (Dai et al., 2020; Gururangan et al.,

2020). For dependency parsing, the use of data from matching genres has been explored by Plank and van Noord (2011), who find improvements for English and Dutch. This is further confirmed for German by Rehbein and Bildhauer (2017).

Automatically inferred topics (Ruder and Plank, 2017) as well as more abstract selection criteria such as overlapping part-of-speech sequences (Søgaard, 2011; Rosa, 2015) have also proven effective at selecting syntactically similar training instances. Vania et al. (2019) further demonstrate that when word embeddings of mutually unintelligible languages align with respect to POS, cross-lingual transfer remains especially effective. With respect to data-driven domain representations, Stymne (2020) shows that treebank embeddings can be used to successfully transfer knowledge from in-domain cross-lingual source treebanks when used in conjunction with in-language, out-of-domain data. In this work, we will rely solely on treebank genre labels as weak supervision and forgo the use of in-language training data as well as instance-level annotations thereof (e.g. POS tags).

Recently, contextualized embeddings have been shown to contain useful information for training data selection. Aharoni and Goldberg (2020) find that clusters formed by embeddings from untuned, monolingual language models correspond well to the genres of their five-domain corpus. Training an English-to-German machine translation model on only the closest embedded sentences to their target 2k-sentence development set outperformed a model trained on the entire dataset.

Although all aforementioned methods assume some degree of in-language training data, our methods will not have access to any annotated target data and will be trained exclusively on out-of-language instances. Building on information stored in pre-trained contextual embeddings, we extend genre-based data selection into the massively multilingual, 104-language, 18-genre setting of Universal Dependencies (Zeman et al., 2020). While previous work further assumed sentence-level genre labels (Ruder and Plank, 2017; Aharoni and Goldberg, 2020), our methods will only have access to treebank-level metadata. An instance’s genre will therefore have to be inferred using weakly supervised approaches. To the best of our knowledge, this constitutes the first application of UD’s instance-level genre distribution to the selection of training data for zero-shot, cross-lingual dependency parsing.

¹Code at <https://personads.me/x/emnlp-2021-code>.




Figure 2: **Genre Distribution in UD.** Ranges indicate upper/lower bounds for sentences per genre inferred from UD metadata. Center marker reflects the distribution under the assumption that genres within treebanks are uniformly distributed. Labels above the bars indicate the number of treebanks which contain each genre.

3 Genre in Universal Dependencies

Universal Dependencies (Nivre et al., 2016) offer annotations for a broad spectrum of languages, with 104 in version 2.7 (Zeman et al., 2020). Of the 1.38 million sentences from the 177 treebanks which we consider, 64 are test-set only and many in this latter third constitute the sole treebank of the language they are in. Such data sparsity becomes even more critical when both the language and the domain are highly specialized and under-resourced.

As more low-resource languages are added in this manner and as the vast majority of the world’s languages remain without annotated data, it becomes important to consider new signals for selecting training data in zero-shot scenarios. If no data in the target language are available, we hypothesize that characteristics of most genres are stable enough across languages to offer a useful guiding criterion for data selection in cross-lingual dependency parsing.

For 26 of the 177 treebanks, their authors have provided sentence-level genre labels. However, these annotations cover only 6% of UD sentences and are typically incompatible across treebanks (with few exceptions such as PUD). At the treebank level, UD fortunately provides 18 approximated genre labels: *academic*, *bible*, *blog*, *email*, *fiction*, *government*, *grammar-examples*, *learner-essays*, *legal*, *medical*, *news*, *nonfiction*, *poetry*, *reviews*, *social*, *spoken*, *web*, *wiki*.

Genres such as *wiki* likely have stronger internal consistency due to cross-lingual creation guidelines. Others such as *fiction* or *web* may have higher variance. While these UD-provided labels are far from perfectly defined (Nivre et al., 2020), they nonethe-

less allow us to operationalize our hypothesis: If genre is globally consistent, it must have a positive effect on cross-lingual transfer performance.

From Figure 2 it is evident that these genres are heavily imbalanced. The minimum number of sentences in a genre is inferred from the sum over the number of instances in treebanks containing only that genre. The upper bound is the sum of all treebanks containing the genre among others. As indicated by these distributional bounds, news articles may constitute up to 70% of the whole UD dataset. Even assuming uniform genre distributions within each treebank (center marker), over half of all sentences in UD would fall into either the *news* or the *non-fiction* category.

Genres with highly specific lexical and/or structural features such as *spoken*, *social* or *medical* are much more underrepresented. Furthermore, they are often only a small part of larger genre mixtures (117 treebanks include multiple genre-labels). These mixtures, with up to 10 genres in one treebank, may contain related genres (e.g. *news*, *non-fiction*, *web*), but also unrelated ones (e.g. *medical*, *poetry*, *social*, *web*) depending on what data was available to authors during annotation.

Out-of-the-box, treebank-level genre labels appear to be highly noisy (see also Nivre et al., 2020). Additionally, individual treebanks are labeled with multiple genres while lacking such labels at the sentence level. We hypothesize that it is therefore necessary to predict *instance-level* genre distributions before targeted data selection can be effective.

4 Targeted Data Selection

In order to measure the effect of genre on the targeted selection of training data, we depart from previous treebank-level selection (Section 2) and introduce three new types of instance-level selection strategies in the following section. They are evaluated on the task of zero-shot dependency parsing in Sections 5 and 6. All of them build on contextualized embeddings learned by the mBERT (Devlin et al., 2019) masked language model (MLM). While MLMs still lack the full breadth of the languages covered in UD (mBERT covers 56 of the 104 languages), they have proven robust in zero-shot scenarios (Devlin et al., 2019; Brown et al., 2020) and have also been found to contain a certain amount of genre information — at least monolingually (Aharoni and Goldberg, 2020; Section 2). We evaluate whether UD’s definition of genre is

also recoverable from these data-driven representations and whether these categories hold cross-lingually.

4.1 Closest Sentence Selection

SENT Akin to the strategy used by [Aharoni and Goldberg \(2020\)](#), this SENTENCE-based method attempts to find the most relevant training data by computing the mean embedding of n unannotated target data samples and retrieving the top- k closest non-target instances according to their cosine distance in embedding space. Notable differences from their original method are the use of a much smaller target data sample ($n = 100$ versus $n = 2000$) as well as the use of mBERT instead of English-only BERT embeddings ([Devlin et al., 2019](#)) due to our cross-lingual setting.

While the monolingual BERT embeddings were found to represent genre to some degree, such MLM embeddings likely contain many more dimensions of semantic and syntactic information. The SENT method alone is therefore not guaranteed to represent data selection by genre as stronger factors may override these signals. Additionally, [Aharoni and Goldberg \(2020\)](#)’s setup assumed five clearly-defined genres with instance-level annotations while UD has 18 genres with varying degrees of specificity which are only defined in the treebank-level metadata.

4.2 Genre Selection

META Separately to MLM embedding-based selection, we evaluate the effectiveness of using the manually assigned genre labels listed in each treebank’s metadata. As seen in Section 3, these labels can be noisy and have variable interpretations across treebanks. Furthermore, each treebank is assigned up to 10 genres, making instance-level selection as in the previous method impossible.

BOOT To bridge this gap to sentence-level selection, we introduce a bootstrapping procedure which iteratively learns an instance-level classifier for UD genre. Each sentence is encoded through mBERT’s CLS token before passing to a classification layer. The model is initialized using standard mBERT weights and begins by training on single-genre treebanks (i.e. standard supervised learning). It then predicts sentence labels for treebanks containing these initial genres. Above a prediction threshold of $0.99 \in [0, 1]$, these are added as new training data for the next round of training. When only

one unclassified genre remains in a treebank, all remaining instances are inferred to be of that last genre. Using this procedure, a single genre label is assigned to each sentence in UD within three steps.

Compared to closest sentence selection (SENT), both of the former methods have the added benefit that no target-data is required in order to make the final training data selection. The training corpus simply consists of all instances labelled as belonging to a genre (BOOT) or to a treebank containing the genre in question (META).

4.3 Closest Cluster Selection

GMM As shown by [Aharoni and Goldberg \(2020\)](#), monolingual BERT embeddings can be clustered into distinct domains using common clustering algorithms such as Gaussian Mixture Models (GMMs). Using mBERT embeddings, we evaluate whether this holds cross-lingually by clustering each treebank into the number of genres which it is said to contain according to the UD-provided metadata. Deviating from previous work, which only uses these clusters for preliminary analyses, we then use them directly for data selection. By computing a mean embedding for each cluster and choosing the closest one to the mean target sample embedding (same as SENT), the most similar data is selected in bulk from each treebank. By only selecting clusters from treebanks for which the metadata states that the target genre is contained, this allows us to identify clusters which most likely correspond to the target genre while avoiding the manual labelling of clusters across 104 languages.

LDA We also evaluate a clustering method based purely on lexical features (i.e. n-grams) instead of pre-trained contextual embeddings. While the selection of the most relevant cluster from each treebank is performed using the same mean embedding distance methodology as for GMM, we use Latent Dirichlet Allocation ([Blei et al., 2001](#); LDA) for the initial clustering step. This decouples the genre-segmentation step from the multitude of non-genre dimensions in the embeddings themselves, while simultaneously not relying on LDA alone for the final data selection (as in [Plank and van Noord, 2011](#); [Mukherjee et al., 2017](#)). Furthermore, this setup allows us to extract genres from languages and scripts unknown to mBERT as well as to compare whether the GMM clusters correspond to those found by using surface-level lexical information alone.

TARGET	AUTHORS	LANGUAGE	FAMILY	MB	SIZE	GENRE
SWL-SSLC	Östling et al. (2017)	Swedish Sign Language	Signed Language	×	203	● spoken
SA-UFAL	Dwivedi and Easha (2017)	Sanskrit	Indo-European	×	230	■ fiction
KPV-Lattice	Partanen et al. (2018)	Komi Zyrian	Uralic	×	435	■ fiction
TA-TTB	Ramasamy and Žabokrtský (2012)	Tamil	Dravidian	✓	600	■ news
GL-TreeGal	Garcia (2016)	Galician	Indo-European	✓	1,000	■ news
YUE-HK	Wong et al. (2017)	Cantonese	Sino-Tibetan	×	1,004	● spoken
CKT-HSE	Tyers and Mishchenkova (2020)	Chukchi	Chukotko-Kamchatkan	×	1,004	● spoken
FO-OFT	Tyers et al. (2018)	Faroese	Indo-European	×	1,208	w wiki
TE-MTG	Rama and Vajjala (2017)	Telugu	Dravidian	✓	1,328	⚡ grammar
MYV-JR	Rueter and Tyers (2018)	Erzya	Uralic	×	1,690	■ fiction
QHE-HIENCS	Bhat et al. (2018)	Hindi-English	Code-Switched	~	1,800	🗨 social
QTD-SAGT	Çetinoğlu and Çöltekin (2019)	Turkish-German	Code-Switched	~	1,891	● spoken

Table 1: **Target Treebanks** with language family (FAMILY), inclusion in mBERT pre-training (MB; included (✓), excluded (×), highly-related languages included (~)), total number of sentences (SIZE) and UD-provided GENRE.

5 Experimental Setup

5.1 Target Treebanks

We evaluate the effect of genre on training data selection using a set of 12 target treebanks from the low-resource end of UD. For our purposes, low-resource is defined as treebanks with more than 200 and less than 2,000 sentences in total and with fewer than 5,000 in-language sentences in UD.

In order to distinguish the effects of genre specifically, we only use single-genre target treebanks and leave the investigation of genre-mixtures to future work. As seen in Table 1, the final set of targets is diverse with respect to genre, language family and their availability during mBERT pre-training.

Only three of the target languages are included in mBERT pre-training, with seven not being covered at all and two having strongly related languages in mBERT’s repertoire: Hindi-English (QHE) → Hindi, English as well as Turkish-German (QTD) → Turkish, German.

The six included genres cover the high-resource *news* (■) and *fiction* (■) as well as the medium resource *wiki* (w) and the lower resource *spoken* (●), *grammar-examples* (⚡) and *social* (🗨).

5.2 Data Selection Setup

In order to train parsers for these largely test-only treebanks, we compare seven proxy training data selection strategies for each target (note that only the first strategy uses in-language training data).

TARGET Where available, we use the true target training split as a performance upper bound against which to compare our methods. These are available for the six targets: SWL-SSLC, TA-TTB, GL-TreeGal, TE-MTG, QHE-HIENCS and QTD-SAGT. For three targets without training splits, we

make use of proxy in-language data: SA-Vedic (Hellwig et al., 2020) for SA-UFAL, KPV-IKDP (Partanen et al., 2018) for KPV-Lattice and FO-FarPaHC (Ingason et al., 2020) for FO-OFT. For the targets YUE-HK, CKT-HSE and MYV-JR no in-language training data are currently available.

RAND selects a random sample of n_{rand} sentences from the non-target-language UD. We do not restrict this selection to treebanks containing the target genre such that data from a more diverse pool of languages may be selected. To ensure an equivalent comparison, we set n_{rand} to the mean of the number of instances selected by BOOT, LDA and GMM (see Appendix C for values of n_{rand}).

SENT selection (see Section 4.1) is based on the mean embedding of 100 target sentences and retrieves the top- k closest out-of-language sentences from all of UD independently of genre. Since k needs to be chosen manually, we set it to the number of instances selected by GMM, which is equally dependent on mBERT embeddings.

META selects all non-target language treebanks which are denoted to contain the target genre (i.e. both single-genre treebanks as well as mixtures). These data pools make up the largest training corpora in our setup (up to 524k instances for *news*) and also subsume the other genre-based selection methods BOOT, LDA and GMM. In this way, it acts as an upper bound in terms of data quantity as well as a baseline for whether treebank-level metadata alone can aid data selection.

BOOT selects only the specific instances classified as being in the target genre for use as training data. The classifier is trained according to the bootstrapping method outlined in Section 4.2. In

SETUP	SWL	SA	KPV	TA	GL	YUE	CKT	FO w	TE	MYV	QHE	QTD	AVG
TARGET	28.01	15.74	13.36	64.05	80.94	—	—	49.55	83.63	—	62.66	55.04	50.28
RAND	3.67	24.81	10.88	50.73	77.65	33.31	15.54	61.88	67.68	20.01	27.01	44.57	36.48
SENT	3.55	23.72	13.71	47.93	77.55	35.78	16.44	62.49	68.05	22.90	26.46	42.74	36.78
META	6.50	24.29	10.22	50.43	76.63	31.19	11.62	61.23	64.91	20.41	9.42	42.58	34.12
BOOT	5.20	21.80	†21.09	49.43	76.66	†49.85	18.40	†66.25	65.56	19.46	14.75	43.80	37.69
GMM	4.85	22.93	†20.91	† 51.53	77.75	† 49.92	† 19.81	†68.25	67.87	20.15	15.09	45.38	38.70
LDA	6.62	23.70	† 22.27	49.17	77.01	†49.40	†19.05	† 68.29	† 68.56	20.54	15.16	44.72	38.71

Table 2: **Zero-shot Parsing Results.** LAS for test splits of target treebanks using training data from target/proxy in-language treebanks (TARGET; where available), random sentence selection (RAND), closest sentence selection (SENT), treebanks containing target genre (META), instances classified as target genre (BOOT) and closest cluster selection (GMM and LDA). Scores marked with † significantly outperform TARGET, RAND, SENT and META.

order to avoid the memorization of target data, we exclude all data in the target languages from the classifier training process.

GMM clusters each treebank into the number of genres denoted by its metadata using mean-pooled mBERT embeddings for each sentence. Training data is then selected according to the closest-cluster procedure outlined in Section 4.3.

LDA works analogously to GMM, but uses LDA to cluster sentences. It uses bags of character 3–6-grams and no language-specific resources (e.g. stop word lists) in order to remain as cross-lingually comparable as possible. Hyperparameters were tuned as outlined in Section 5.3.

All methods relying on unannotated target data for the data selection process use 100 random sentences from the target treebank (changes across random initializations). In practical terms, this corresponds to having access to a small amount of target-like data — without gold dependency structures — and selecting the best possible training data for which we do have annotations.

Alternatively, BOOT (as well as META and RAND implicitly) work in a fully zero-shot manner as we only assume knowledge of the intended target genre, but do not assume access to the target sentences *nor* their annotations.

5.3 Training Setup

We use the biaffine attention parser (Dozat and Manning, 2017) implementation of MaChAmp v0.2 (van der Goot et al., 2021) with default hyperparameters. Each step involving non-deterministic components is rerun using three random seeds.

For efficiency reasons, the seven largest treebanks were subsampled to 20k instances per split.

Performance is measured using the labeled attachment scores (LAS) averaged across random initializations. Additionally, we report unlabeled attachment scores (UAS), the number of selected instances as well as the variance across runs in Appendix C. Significance is evaluated at $\alpha < 0.05$ using a paired bootstrapped sign test with 10k resampling and Bonferroni correction (Bonferroni, 1936) for the multiple comparisons across random initializations. Appendix B lists all additional hyperparameter settings.

It is important to note that besides the upper bound in-language setup (TARGET), no parser is trained on in-language data. For the tuning of method-specific hyperparameters (LDA features, BOOT thresholds), development sets of the five treebanks containing such splits were used: SWL-SSLC, TA-TTB, TE-MTG, QHE-HIENCS and QTD-SAGT (details in Appendix B). During parser training, development data for early stopping is based solely on the out-of-language data selected by each method and not on the in-language target data itself (also excluding constituent languages for code switched targets). Results are reported on each target’s test set without any further tuning.

6 Results

6.1 Zero-shot Parsing Results

As expected, Table 2 shows that training the parser on target data (TARGET) results in the best overall performance even though the training corpora for these setups almost never exceed 1k instances. The target treebanks for which in-language data are available, consolidate into a final average of 50.28 LAS. This highlights the overall difficulty of parsing these low-resource treebanks. As the parser is initialized using mBERT, the scores on Tamil (TA),

Galician (GL) and Telugu (TE), which are included in its pre-training, are highest overall compared to non-included languages or code-switched variants.

It is noteworthy that when a same-language proxy treebank was used for parser training, scores are lower compared to the other methods. In these three cases, namely Sanskrit (SA), Komi Zyrian (KPV) and Faroese (FO), none of the proxy treebanks include the target’s genre which may be a strong contributing factor to this discrepancy.

Turning to our zero-shot setups, META data selection based on treebank-level annotations alone performs worst overall at 34.12 LAS despite constituting the largest training corpora in each setup (see Appendix C for details). Compared to the TARGET upper bounds, it shows how training on two orders of magnitudes more data can still be insufficient if they do not follow the target distribution.

Both RAND and SENT outperform the META baseline at 36.48 and 36.78 LAS respectively. These aggregated scores also highlight that sentence-based selection alone insufficiently captures cross-lingual characteristics as to outperform random chance in most cases.

In contrast, combining latent information in the MLM embeddings with higher-level genre information leads to performance increases not achievable by either method alone. Both GMM and LDA achieve the highest scores across the majority of target treebanks and the highest cross-lingual averages of 38.70 LAS and 38.71 LAS respectively. These scores reflect their similar performance across targets, however we do observe that LDA achieves slightly higher scores on languages which are not included in mBERT pre-training: e.g. Swedish Sign Language (SWL), Sanskrit (SA) and Komi Zyrian (KPV). We hypothesize that this is a result of GMM’s dependence on latent information in the mBERT embeddings while LDA constructs clusters independently, based solely on surface-level lexical features (i.e. n-grams).

Finally, amplifying genre information in the mBERT embeddings using our BOOT method also leads to performance increases compared to using untuned embeddings or the coarser grained treebank-level metadata. While it does not entirely reach the performance of the cluster selection methods, its overall average of 37.69 LAS as well as generally similar performance patterns to LDA and GMM lead us to believe that all three methods are picking up on and are amplifying similar latent

genre information. As an added benefit, BOOT is able to reach this competitive performance without the need for any target data samples (as opposed to GMM and LDA which use 100 raw samples for cluster selection).

Using our proposed genre-based selection methods we are therefore able to consistently outperform in-language/out-of-genre upper bounds for these low-resource target treebanks. Comparing our results to van der Goot et al. (2021) who train an identical parser architecture on each UD treebank’s respective training split, proxy treebank (for test-only) or all of UD, our methods significantly outperform their best models on five of twelve target treebanks.² There are significant increases for both SA-UFAL (16.5 → 23.7 LAS) and KPV-Lattice (11.7 → 22.3 LAS).³ For the targets YUE-HK (32.7 → 49.9 LAS), CKT-HSE (15.3 → 19.8 LAS) and FO-OFT (62.7 → 68.3 LAS), these scores furthermore constitute — to the best of our knowledge — state-of-the-art results without requiring annotated in-language data.

6.2 Analysis of Selected Data

Further analyzing the patterns of data selection allows us to identify some of the reasons behind the differences in performance (visualizations can be found in Appendix D).

RAND closely follows the overall data distribution in UD, selecting the most instances from the largest treebanks such as German-HDT (Borges Völker et al., 2019) and selecting none to almost none from low-resource treebanks. SENT follows a similar distribution albeit rarely selecting zero instances from any given language. This behaviour does not change substantially between targets, indicating less targeted data selection.

While the larger language diversity of the aforementioned RAND and SENT does not seem to be enough to outperform genre-selection in most cases, it can be helpful when in-genre data is not as linguistically diverse. For the targets SA-UFAL and MYV-JR (*fiction*) both methods outperform genre-based selection by around 2% LAS.

A clear example of insufficient in-genre data is the QHE-HIENCS target. It represents a highly-specialized variation of the *social* genre, specifically Twitter data. Although the genre-based selec-

²We compare against the highest score across all of their proposed models for each treebank.

³Dehouck and Denis (2019) achieve higher scores using a parsing architecture with POS and morphological features.

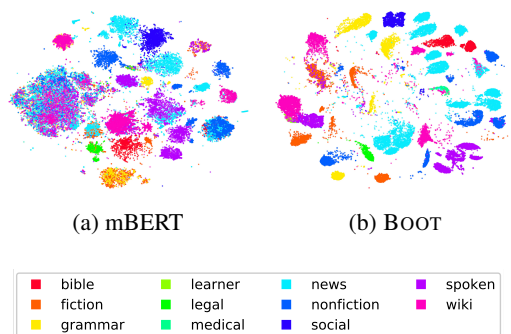


Figure 3: **UD Genres in Embedding Space** of (a) untuned mBERT and (b) genre-tuned BOOT. Sentences from single-genre treebanks (up to 1k each) colored by genre, plotted using tSNE (Maaten and Hinton, 2008). Tuning using genre as weak supervision clearly amplifies genre information.

tion methods correctly identify and cluster the Italian Twitter data from IT-PoSTWITA (Sanguinetti et al., 2018) and IT-TWITTIRO (Cignarella et al., 2019), there is a lack of such in-genre data from other languages,⁴ leading these parsers to overfit on Italian specifically. This once again highlights the difficulty of selecting proxy training data which covers all desired characteristics — even from a dataset as diverse as UD.

In general, the genre-driven methods make fairly similar selections given their shared baseline pool of treebanks containing the target genre in-mixture (see Appendix D). Since using all of these data however results in the worst overall performance (META) while BOOT, GMM and LDA perform best, the targeted selection of relevant subsets within the larger META pool appears to be key. Frequently, large treebanks such as Polish-LFG (Patejuk and Przepiórkowski, 2018) with 14k instances from *fiction*, *news*, *nonfiction*, *social* and *spoken* are subsampled to a much smaller fraction (around 3k instances in this example). The fact that these proportions as well as the selected instances themselves are relatively consistent across same-genre targets corroborates that all our methods are picking up on similar, data-driven notions of genre.

Figure 3 further visualizes the presence of latent genre using t-SNE plots of up to 1k randomly sampled sentence embeddings from each of UD’s single-genre treebanks. In their untuned state (Figure 3a), some local genre clusters do manifest.

⁴More non-official Twitter-based treebanks in UD style exist (Sanguinetti et al., 2020) which were left out of this study as they are not part of UD and contain annotation divergences.

However, these mainly correspond to specialized treebanks such as the aforementioned Italian Twitter treebanks (*social*). Most other genres occur in language-level mixtures or in a large overall “blob” on the left. By amplifying genre explicitly using the BOOT procedure, each individual genre is much more clearly segmented (Figure 3b).

In conclusion, the presence of similar performance patterns across all our proposed genre-driven methods — while having separate approaches to treebank segmentation (weakly supervised tuning for BOOT, treebank-internal embedding distances for GMM, n-grams for LDA) — confirms our hypothesis that instance-level genre can be identified cross-lingually from contextualized representations and aids zero-shot parsing.

7 Conclusions

In absence of in-language training data, we have explored UD-specified genre as an alternative signal for data selection. While prior work had indicated the presence of genre information in monolingual contextualized embeddings (Aharoni and Goldberg, 2020), an analogous strategy using mBERT embeddings proved insufficient in the cross-lingual parsing setting (SEN), performing close to the random baseline (RAND). Relying on manual, treebank-level genre labels (META) proved even less performant, producing the lowest scores despite corresponding to a practitioner’s typical first choice of selecting the largest number of training instances.

In order to enable finer-grained, instance-level data selection, we proposed three methods for combining latent genre information in the unsupervised contextualized representations with the treebank metadata: weakly supervised BOOT, sentence embedding-based GMM and n-gram-based LDA. Despite their different approaches to treebank segmentation, each method significantly outperformed the purely embedding-based SENT as well as the metadata (META) and random baselines (RAND). Their similar performance patterns and selected data distributions further indicate that each method is identifying a shared, data-driven notion of genre.

For future work, it will be important to extend our proposed approaches beyond single-genre targets towards genre-mixtures and more treebanks overall. As the data selected by these methods is further limited by the number of treebanks in each respective genre, combining a larger set of selection signals will be equally crucial.

Acknowledgments

We would like to thank the NLPnorth group for insightful discussions on this work — in particular Elisa Bassignana, Maria Barrett and Mike Zhang. Thanks to Héctor Martínez Alonso for feedback on an early draft as well as ITU’s High-performance Computing Cluster team. Finally, we thank the anonymous reviewers for their thorough feedback. This research is supported by the Independent Research Fund Denmark (DFF) grant 9063-00077B and an Amazon Faculty Research (ARA) Award.

References

- Roei Aharoni and Yoav Goldberg. 2020. [Unsupervised domain clusters in pretrained language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7747–7763, Online. Association for Computational Linguistics.
- Irshad Bhat, Riyaz A. Bhat, Manish Shrivastava, and Dipti Sharma. 2018. [Universal Dependency parsing for Hindi-English code-switching](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 987–998, New Orleans, Louisiana. Association for Computational Linguistics.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2001. [Latent dirichlet allocation](#). In *Advances in Neural Information Processing Systems 14 [Neural Information Processing Systems: Natural and Synthetic, NIPS 2001, December 3-8, 2001, Vancouver, British Columbia, Canada]*, pages 601–608. MIT Press.
- Carlo Bonferroni. 1936. Teoria statistica delle classi e calcolo delle probabilità. *Pubblicazioni del R Istituto Superiore di Scienze Economiche e Commerciali di Firenze*, 8:3–62.
- Emanuel Borges Völker, Maximilian Wendt, Felix Henning, and Arne Köhn. 2019. [HDT-UD: A very large Universal Dependencies treebank for German](#). In *Proceedings of the Third Workshop on Universal Dependencies (UDW, SyntaxFest 2019)*, pages 46–57, Paris, France. Association for Computational Linguistics.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Özlem Çetinoğlu and Çağrı Çöltekin. 2019. [Challenges of annotating a code-switching treebank](#). In *Proceedings of the 18th International Workshop on Treebanks and Linguistic Theories (TLT, SyntaxFest 2019)*, pages 82–90, Paris, France. Association for Computational Linguistics.
- Alessandra Teresa Cignarella, Cristina Bosco, and Paolo Rosso. 2019. [Presenting TWITTIRÒ-UD: An Italian Twitter treebank in Universal Dependencies](#). In *Proceedings of the Fifth International Conference on Dependency Linguistics (Depling, SyntaxFest 2019)*, pages 190–197, Paris, France. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Xiang Dai, Sarvnaz Karimi, Ben Hachey, and Cecile Paris. 2020. [Cost-effective selection of pretraining data: A case study of pretraining BERT on social media](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1675–1681, Online. Association for Computational Linguistics.
- Mathieu Dehouck and Pascal Denis. 2019. [Phylogenetic multi-lingual dependency parsing](#). 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 192–203, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Timothy Dozat and Christopher D. Manning. 2017. [Deep biaffine attention for neural dependency parsing](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.

- Puneet Dwivedi and Guha Easha. 2017. Universal Dependencies for Sanskrit. *International Journal of Advance Research, Ideas and Innovations in Technology*, 3(4).
- Fonticons. 2021. [Font Awesome Icons](#). CC-BY 4.0 License.
- Marcos Garcia. 2016. Universal dependencies guidelines for the Galician-TreeGal treebank. Technical report, Technical Report, LyS Group, Universidade da Coruna.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don't stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Oliver Hellwig, Salvatore Scarlata, Elia Ackermann, and Paul Widmer. 2020. [The treebank of vedic Sanskrit](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5137–5146, Marseille, France. European Language Resources Association.
- Anton Karl Ingason, Eiríkur Rögnvaldsson, Einar Freyr Sigurosson, and Joel C. Wallenberg. 2020. [The Faroese parsed historical corpus](#). CLARIN-IS, Stofnun Árna Magnússonar.
- Brett Kessler, Geoffrey Nunberg, and Hinrich Schütze. 1997. [Automatic detection of text genre](#). In *35th Annual Meeting of the Association for Computational Linguistics and 8th Conference of the European Chapter of the Association for Computational Linguistics*, pages 32–38, Madrid, Spain. Association for Computational Linguistics.
- Dan Kondratyuk and Milan Straka. 2019. [75 languages, 1 model: Parsing Universal Dependencies universally](#). 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 2779–2795, Hong Kong, China. Association for Computational Linguistics.
- Anne Lauscher, Vinit Ravishankar, Ivan Vulić, and Goran Glavaš. 2020. [From zero to hero: On the limitations of zero-shot language transfer with multilingual Transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4483–4499, Online. Association for Computational Linguistics.
- David Y. W. Lee. 2001. [Genres, registers, text types, domains and styles: Clarifying the concepts and navigating a path through the BNC jungle](#). *Language Learning & Technology*, 5:37–72.
- Ilya Loshchilov and Frank Hutter. 2017. [Fixing weight decay regularization in adam](#). *Computing Research Repository*, arXiv: 1711.05101. Version 3.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605.
- Ailsa Meechan-Maddon and Joakim Nivre. 2019. [How to parse low-resource languages: Cross-lingual parsing, target language annotation, or both?](#) In *Proceedings of the Fifth International Conference on Dependency Linguistics (Depling, SyntaxFest 2019)*, pages 112–120, Paris, France. Association for Computational Linguistics.
- Atreyee Mukherjee, Sandra Kübler, and Matthias Scheutz. 2017. [Creating POS tagging and dependency parsing experts via topic modeling](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 347–355, Valencia, Spain. Association for Computational Linguistics.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajič, Christopher D. Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. [Universal Dependencies v1: A multilingual treebank collection](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1659–1666, Portorož, Slovenia. European Language Resources Association (ELRA).
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher D. Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. [Universal Dependencies v2: An evergrowing multilingual treebank collection](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4034–4043, Marseille, France. European Language Resources Association.
- Robert Östling, Carl Börstell, Moa Gärdénfors, and Mats Wirén. 2017. [Universal Dependencies for Swedish Sign Language](#). In *Proceedings of the 21st Nordic Conference on Computational Linguistics*, pages 303–308, Gothenburg, Sweden. Association for Computational Linguistics.
- Niko Partanen, Rogier Blokland, KyungTae Lim, Thierry Poibeau, and Michael Rießler. 2018. [The first Komi-Zyrian Universal Dependencies treebanks](#). In *Proceedings of the Second Workshop on Universal Dependencies (UDW 2018)*, pages 126–132, Brussels, Belgium. Association for Computational Linguistics.
- Agnieszka Patejuk and Adam Przepiórkowski. 2018. [From Lexical Functional Grammar to Enhanced Universal Dependencies: Linguistically informed treebanks of Polish](#). Institute of Computer Science, Polish Academy of Sciences, Warsaw.

- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. [Scikit-learn: Machine learning in Python](#). *Journal of Machine Learning Research*, 12:2825–2830.
- Barbara Plank. 2011. [Domain adaptation for parsing](#). University of Groningen.
- Barbara Plank and Gertjan van Noord. 2011. [Effective measures of domain similarity for parsing](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1566–1576, Portland, Oregon, USA. Association for Computational Linguistics.
- Taraka Rama and Sowmya Vajjala. 2017. [A Telugu treebank based on a grammar book](#). In *Proceedings of the 16th International Workshop on Treebanks and Linguistic Theories*, pages 119–128, Prague, Czech Republic.
- Loganathan Ramasamy and Zdeněk Žabokrtský. 2012. [Prague dependency style treebank for Tamil](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 1888–1894, Istanbul, Turkey. European Language Resources Association (ELRA).
- Ines Rehbein and Felix Bildhauer. 2017. [Data point selection for genre-aware parsing](#). In *Proceedings of the 16th International Workshop on Treebanks and Linguistic Theories*, pages 95–105, Prague, Czech Republic.
- Rudolf Rosa. 2015. [Parsing natural language sentences by semi-supervised methods](#). *Computing Research Repository*, arXiv: 1506.04897. Version 1.
- Sebastian Ruder and Barbara Plank. 2017. [Learning to select data for transfer learning with Bayesian optimization](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 372–382, Copenhagen, Denmark. Association for Computational Linguistics.
- Jack Rueter and Francis Tyers. 2018. [Towards an open-source universal-dependency treebank for Erzya](#). In *Proceedings of the Fourth International Workshop on Computational Linguistics of Uralic Languages*, pages 106–118, Helsinki, Finland. Association for Computational Linguistics.
- Manuela Sanguinetti, Cristina Bosco, Lauren Cassidy, Özlem Çetinoğlu, Alessandra Teresa Cignarella, Teresa Lynn, Ines Rehbein, Josef Ruppenhofer, Djamé Seddah, and Amir Zeldes. 2020. [Treebanking user-generated content: A proposal for a unified representation in Universal Dependencies](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5240–5250, Marseille, France. European Language Resources Association.
- Manuela Sanguinetti, Cristina Bosco, Alberto Lavelli, Alessandro Mazzei, Oronzo Antonelli, and Fabio Tamburini. 2018. [PoSTWITA-UD: an Italian Twitter treebank in Universal Dependencies](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Anders Søgaard. 2011. [Data point selection for cross-language adaptation of dependency parsers](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 682–686, Portland, Oregon, USA. Association for Computational Linguistics.
- Sara Stymne. 2020. [Cross-lingual domain adaptation for dependency parsing](#). In *Proceedings of the 19th International Workshop on Treebanks and Linguistic Theories*, pages 62–69, Düsseldorf, Germany. Association for Computational Linguistics.
- Francis Tyers and Karina Mishchenkova. 2020. [Dependency annotation of noun incorporation in polysynthetic languages](#). In *Proceedings of the Fourth Workshop on Universal Dependencies (UDW 2020)*, pages 195–204, Barcelona, Spain (Online). Association for Computational Linguistics.
- Francis Tyers, Mariya Sheyanova, Aleksandra Martynova, Pavel Stepachev, and Konstantin Vinogorodskiy. 2018. [Multi-source synthetic treebank creation for improved cross-lingual dependency parsing](#). In *Proceedings of the Second Workshop on Universal Dependencies (UDW 2018)*, pages 144–150, Brussels, Belgium. Association for Computational Linguistics.
- Ahmet Üstün, Arianna Bisazza, Gosse Bouma, and Gertjan van Noord. 2020. [UDapter: Language adaptation for truly Universal Dependency parsing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2302–2315, Online. Association for Computational Linguistics.
- Rob van der Goot, Ahmet Üstün, Alan Ramponi, Ibrahim Sharaf, and Barbara Plank. 2021. [Massive choice, ample tasks \(MaChAmp\): A toolkit for multi-task learning in NLP](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 176–197, Online. Association for Computational Linguistics.
- 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.

- Bonnie Webber. 2009. [Genre distinctions for discourse in the Penn TreeBank](#). In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 674–682, Suntec, Singapore. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Tak-sum Wong, Kim Gerdes, Herman Leung, and John Lee. 2017. [Quantitative comparative syntax on the Cantonese-Mandarin parallel dependency treebank](#). In *Proceedings of the Fourth International Conference on Dependency Linguistics (Depling 2017)*, pages 266–275, Pisa, Italy. Linköping University Electronic Press.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#). *Computing Research Repository*, arXiv: 1609.08144. Version 2.
- Daniel Zeman, Joakim Nivre, Mitchell Abrams, Elia Ackermann, Noëmi Aeppli, Hamid Aghaei, Željko Agić, Amir Ahmadi, Lars Ahrenberg, Chika Kennedy Ajede, Gabrielė Aleksandravičiūtė, Ika Alfina, Lene Antonsen, Katya Aplonova, Angelina Aquino, Carolina Aragon, Maria Jesus Aranzabe, Hórunn Arnardóttir, Gashaw Arutie, Jessica Naraiswari Arwidarasti, Masayuki Asahara, Luma Ateyah, Furkan Atmaca, Mohammed Attia, Aitziber Atutxa, Liesbeth Augustinus, Elena Badmaeva, Keerthana Balasubramani, Miguel Ballesteros, Esha Banerjee, Sebastian Bank, Verginica Barbu Mititelu, Victoria Basmov, Colin Batchelor, John Bauer, Seyyit Talha Bedir, Kepa Bengoetxea, Gözde Berk, Yevgeni Berzak, Irshad Ahmad Bhat, Riyaz Ahmad Bhat, Erica Biagetti, Eckhard Bick, Agnė Bielinskienė, Kristín Bjarnadóttir, Rogier Blokland, Victoria Bobicev, Loïc Boizou, Emanuel Borges Völker, Carl Börstell, Cristina Bosco, Gosse Bouma, Sam Bowman, Adriane Boyd, Kristina Brokaitė, Aljoscha Burchardt, Marie Candito, Bernard Caron, Gauthier Caron, Tatiana Cavalcanti, Gülşen Cebiroğlu Eryiğit, Flavio Massimiliano Cecchini, Giuseppe G. A. Celano, Slavomír Čéplö, Savas Cetin, Özlem Çetinoğlu, Fabricio Chalub, Ethan Chi, Yongseok Cho, Jinho Choi, Jayeol Chun, Alessandra T. Cignarella, Silvie Cinková, Aurélie Collomb, Çağrı Çöltekin, Miriam Connor, Marine Courtin, Elizabeth Davidson, Marie-Catherine de Marneffe, Valeria de Paiva, Mehmet Oguz Derin, Elvis de Souza, Arantza Diaz de Ilaraza, Carly Dickerson, Arawinda Dinakaramani, Bamba Dione, Peter Dirix, Kaja Dobrovoljc, Timothy Dozat, Kira Droganova, Puneet Dwivedi, Hanne Eckhoff, Marhaba Eli, Ali Elkahky, Binyam Ephrem, Olga Erina, Tomaž Erjavec, Aline Etienne, Wograin Evelyn, Sidney Facundes, Richárd Farkas, Marília Fernanda, Hector Fernandez Alcalde, Jennifer Foster, Cláudia Freitas, Kazunori Fujita, Katarína Gajdošová, Daniel Galbraith, Marcos Garcia, Moa Gärdenfors, Sebastian Garza, Fabrício Ferraz Gerardi, Kim Gerdes, Filip Ginter, Iakes Goenaga, Koldo Gojenola, Memduh Gökırmak, Yoav Goldberg, Xavier Gómez Guinovart, Berta González Saavedra, Bernadeta Griciūtė, Matias Gioni, Loïc Grobol, Normunds Grūzītis, Bruno Guillaume, Céline Guillot-Barbance, Tunga Güngör, Nizar Habash, Hinrik Hafsteinnsson, Jan Hajič, Jan Hajič jr., Mika Hämäläinen, Linh Hà Mỹ, Na-Rae Han, Muhammad Yudistira Hanifmuti, Sam Hardwick, Kim Harris, Dag Haug, Johannes Heinecke, Oliver Hellwig, Felix Henning, Barbora Hladká, Jaroslava Hlaváčová, Florinel Hociung, Petter Hohle, Eva Huber, Jena Hwang, Takumi Ikeda, Anton Karl Ingason, Radu Ion, Elena Irimia, Oľájdé Ishola, Tomáš Jelínek, Anders Johannsen, Hildur Jónsdóttir, Fredrik Jørgensen, Markus Juutinen, Sarveswaran K, Hüner Kaşıkara, Andre Kaasen, Nadezhda Kabaeva, Sylvain Kahane, Hiroshi Kanayama, Jenna Kanerva, Boris Katz, Tolga Kayadelen, Jessica Kenney, Václava Kettnerová, Jesse Kirchner, Elena Klementieva, Arne Köhn, Abdullatif Köksal, Kamil Kopacewicz, Timo Korhakangas, Natalia Kotsyba, Jolanta Kovalevskaitė, Simon Krek, Parameswari Krishnamurthy, Sookyoung Kwak, Veronika Laippala, Lucia Lam, Lorenzo Lambertino, Tatiana Lando, Septina Dian Larasati, Alexei Lavrentiev, John Lee, Phuong Lê Hồng, Alessandro Lenci, Saran Lertpradit, Herman Leung, Maria Levina, Cheuk Ying Li, Josie Li, Keying Li, Yuan Li, KyungTae Lim, Krister Lindén, Nikola Ljubešić, Olga Loginova, Andry Luthfi, Mikko Luukko, Olga Lyashevskaya, Teresa Lynn, Vivien Macketanz, Aibek Makazhanov, Michael Mandl, Christopher Manning, Ruli Manurung, Cătălina Mărânduc, David Mareček, Katrin Marheinecke, Héctor Martínez Alonso, André Martins, Jan Mašek, Hiroshi Matsuda, Yuji Matsumoto, Ryan McDonald, Sarah McGuinness, Gustavo Mendonça, Niko Miekka, Karina Mischenkova, Margarita Misirpashayeva, Anna Missilä, Cătălin Mititelu, Maria Mitrofan, Yusuke Miyao, AmirHossein Mojiri Foroushani, Amirsaeid Moloodi, Simonetta

Montemagni, Amir More, Laura Moreno Romero, Keiko Sophie Mori, Shinsuke Mori, Tomohiko Morioka, Shigeki Moro, Bjartur Mortensen, Bohdan Moskalevskyi, Kadri Muischnek, Robert Munro, Yugo Murawaki, Kaili Müürisep, Pinkey Nainwani, Mariam Nakhlé, Juan Ignacio Navarro Horňáček, Anna Nedoluzhko, Gunta Nešpore-Bērzkalne, Luong Nguyễn Thị, Huyền Nguyễn Thị Minh, Yoshihiro Nikaido, Vitaly Nikolaev, Rattima Nitisaroj, Alireza Nourian, Hanna Nurmi, Stina Ojala, Atul Kr. Ojha, Adédayo Olúòkun, Mai Omura, Emeka Onwuegbuzia, Petya Osenova, Robert Östling, Lilja Óvrelid, Şaziye Betül Özateş, Arzucan Özgür, Balkız Öztürk Başaran, Niko Partanen, Elena Pascual, Marco Passarotti, Agnieszka Patejuk, Guilherme Paulino-Passos, Angelika Peljak-Łapińska, Siyao Peng, Cenel-Augusto Perez, Natalia Perkova, Guy Perrier, Slav Petrov, Daria Petrova, Jason Phelan, Jussi Piitulainen, Tommi A Pirinen, Emily Pitler, Barbara Plank, Thierry Poibeau, Larisa Ponomareva, Martin Popel, Lauma Pretkalniņa, Sophie Prévost, Prokopis Prokopidis, Adam Przepiórkowski, Tiina Puolakainen, Sampo Pyysalo, Peng Qi, Andriela Rääbis, Alexandre Rademaker, Taraka Rama, Loganathan Ramasamy, Carlos Ramisch, Fam Rashel, Mohammad Sadegh Rasooli, Vinit Ravishankar, Livy Real, Petru Rebeja, Siva Reddy, Georg Rehm, Ivan Riabov, Michael Rießler, Erika Rimkutė, Larissa Rinaldi, Laura Rituma, Luisa Rocha, Eiríkur Rögnvaldsson, Mykhailo Romanenko, Rudolf Rosa, Valentin Roşca, Davide Rovati, Olga Rudina, Jack Rueter, Kristján Rúnarsson, Shoal Sadde, Pegah Safari, Benoît Sagot, Aleksí Sahala, Shadi Saleh, Alessio Salomoni, Tanja Samardžić, Stephanie Samson, Manuela Sanguinetti, Dage Särg, Baiba Saulīte, Yanin Sawanakunanon, Kevin Scannell, Salvatore Scarlata, Nathan Schneider, Sebastian Schuster, Djámé Seddah, Wolfgang Seeker, Mojgan Seraji, Mo Shen, Atsuko Shimada, Hiroyuki Shirasu, Muh Shohibussirri, Dmitry Sichinava, Einar Freyr Sigurðsson, Aline Silveira, Natalia Silveira, Maria Simi, Radu Simionescu, Katalin Simkó, Mária Šimková, Kiril Simov, Maria Skachedubova, Aaron Smith, Isabela Soares-Bastos, Carolyn Spadine, Steinhór Steingrímsson, Antonio Stella, Milan Straka, Emmett Strickland, Jana Strnadová, Alane Suhr, Yogi Lesmana Sulestio, Umut Sulubacak, Shingo Suzuki, Zsolt Szántó, Dima Taji, Yuta Takahashi, Fabio Tamburini, Mary Ann C. Tan, Takaaki Tanaka, Samson Tella, Isabelle Tellier, Guillaume Thomas, Lisi Torga, Marsida Toska, Trond Trosterud, Anna Trukhina, Reut Tsarfaty, Utku Türk, Francis Tyers, Sumire Uematsu, Roman Untilov, Zdeňka Urešová, Larraitz Uria, Hans Uszkoreit, Andrius Utká, Sowmya Vajjala, Daniel van Niekerk, Gertjan van Noord, Viktor Varga, Eric Villemonte de la Clergerie, Veronika Vincze, Aya Wakasa, Joel C. Wallenberg, Lars Wallin, Abigail Walsh, Jing Xian Wang, Jonathan North Washington, Maximilian Wendt, Paul Widmer, Seyi Williams, Mats Wirén, Christian Wittern, Tsegay Woldemariam, Tak-sum Wong, Alina Wróblewska, Mary Yako, Kayo Yamashita, Naoki Yamazaki, Chunxiao Yan, Koichi Yasuoka,

Marat M. Yavrumyan, Zhuoran Yu, Zdeněk Žabokrtský, Shorouq Zahra, Amir Zeldes, Hanzhi Zhu, and Anna Zhuravleva. 2020. [Universal dependencies 2.7](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.

Appendix

A Universal Dependencies Setup

All experiments make use of Universal Dependencies v2.7 (Zeman et al., 2020; UD). From the total set of 183 treebanks, we use all except for the following six (due to licensing restrictions): *AR-NYUAD*, *EN-ESL*, *EN-GUMReddit*, *FR-FTB*, *JA-BCCWJ*, *GUN-Dooley*. In total 1.38 million sentences are used in our experiments.

Target Treebanks As listed in the main paper, our target treebanks are *Swedish Sign Language-SSLC* (Östling et al., 2017), *Sanskrit-UFAL* (Dwivedi and Easha, 2017), *Komi Zyrian-Lattice* (Partanen et al., 2018), *Tamil-TTB* (Ramasamy and Žabokrtský, 2012), *Galician-TreeGal* (Garcia, 2016), *Cantonese-HK* (Wong et al., 2017), *Chukchi-HSE* (Tyers and Mishchenkova, 2020), *Faroese-OFT* (Tyers et al., 2018), *Telugu-MTG* (Rama and Vajjala, 2017), *Erzya-JR* (Rueter and Tyers, 2018), *Hindi-English-HIENCS* (Bhat et al., 2018) and *Turkish-German-SAGT* (Çetinoğlu and Çöltekin, 2019).

Development Data For the initial tuning of LDA input features as well as the bootstrapping threshold, we used the only five treebanks with development data: *SWL-SSLC*, *TA-TTB*, *TE-MTG*, *QHE-HIENCS*, *QTD-SAGT*.

For the early stopping of parser training, no such in-language validation data is used (to ensure a pure zero-shot setup). Instead, the data selected by each selection method is split in an 80%/20% fashion and is used as a proxy, out-of-language development split.

Similarly, the training of the bootstrapping classifier (BOOT) uses only the non-target-language portion of UD (i.e. excluding all treebanks of the 12 target languages plus constituent languages for code-switched). For efficiency reasons, this data is further subsampled to 40k total instances. Both the training and validation (used for early stopping) of BOOT are therefore similarly conducted without any target-language data.

Subsets Since data selection is at the core of this research, the exact instance IDs of each subset are available in the supplementary code.

B Model and Training Details

The following describes architecture and training details for all methods. When not further defined,

default hyperparameters are used. Implementations are available in the supplementary code.

Infrastructure Neural models are trained on an NVIDIA A100 GPU with 40 GB of VRAM. Since most of our experiments do not require MLM sentence embeddings to be updated, we compute them once and store them on disk to save GPU cycles.

Multilingual Language Model The MLM used in this work is mBERT (Devlin et al., 2019) as implemented in the Transformers library (Wolf et al., 2020)⁵. Embeddings are of size $d_{\text{emb}} = 768$ and the model itself has 178 million total parameters. To create sentence embeddings in the SENT and GMM methods, we use the mean-pooled WordPiece embeddings (Wu et al., 2016) of the final layer.

Clustering Methods Both *Gaussian Mixture Models* (GMM) and *Latent Dirichlet Allocation* (Blei et al., 2001; LDA) use implementations from scikit-learn v0.23 (Pedregosa et al., 2011). LDA uses bags of character 3–6-grams which occur in at least two and in at most 30% of sentences. The n-gram sizes were initially tuned on target treebanks with available development sets (see Appendix A). We found character 1–5-grams to perform approximately 2.5 LAS worse and word unigrams to perform approximately 2 LAS worse than the final method. GMMs use the mBERT sentence embeddings directly as input. Both methods are CPU-bound and complete the clustering of all treebanks in UD in under 45 minutes.

Bootstrapping (BOOT) builds on the standard mBERT architecture as follows: mBERT \rightarrow CLS \rightarrow linear layer ($d_{\text{emb}} \times 18$) \rightarrow softmax. The training has an epoch limit of 100 with early stopping after 3 iterations without improvements on the development set. No target-language data is used during this process. An alternate bootstrapping threshold of 0.9 was evaluated and found to perform approximately 1 LAS worse on the development subset (see Appendix A) than the final value of 0.99. Backpropagation is performed using AdamW (Loshchilov and Hutter, 2017) with a learning rate of 10^{-7} on batches of size 16. The fine-tuning procedure requires GPU hardware which can host mBERT, corresponding to 10 GB of VRAM. Training on the subsampled 40k instance, non-target-language data takes approximately seven hours.

⁵bert-base-multilingual-cased

Dependency Parsers Every parsing experiment in the main paper uses a biaffine attention parser (Dozat and Manning, 2017) implemented in the MaChAmp v0.2 framework (van der Goot et al., 2021) using default hyperparameters. The sentence encoder is initialized with standard mBERT weights. The training duration is foremost dependent on input data quantity. For the largest corpus (META for TA-TTB with 524k instances) this corresponds to 55 hours. Our proposed methods create smaller, targeted training corpora (around 80k instances on average) such that a better performing parser can be trained in approximately 90 minutes on the same hardware.

Random Initializations Each experiment is run thrice using the seeds 41, 42 and 43. This relates to the random subsampling of data as well as to model initialization (both parsers and selection).

C Additional Results

In addition to the labeled attachment scores (LAS) reported in the main paper, we list LAS standard deviation across random initializations in Table 5, unlabeled attachment scores (UAS) in Table 4 as well as the number of selected training instances per method in Table 3.

Predictions We additionally provide the instance-level predictions of each method and each random initialization as CoNLL-U files in the supplementary material in order ensure that future work can evaluate the statistical significance of performance differences.

D Data Selection Analysis

Figure 4 displays the distribution of selected instances across all treebanks of UD per target treebank and method. Proportions are normalized to $[0, 1]$ for each method (i.e. across each column). Due to the large number of cells, we recommend viewing this figure digitally.

SETUP	SWL ●	SA ■	KPV ■	TA ■	GL ■	YUE ●	CKT ●	FO w	TE ✗	MYV ■	QHE ✎	QTD ●	AVG
TARGET	87	3k	132	400	600	—	—	1k	1k	—	1k	285	839
RAND	31k	81k	84k	249k	244k	30k	30k	50k	21k	86k	12k	30k	79k
SENT	33k	95k	101k	271k	236k	31k	30k	58k	23k	113k	14k	31k	86k
META	62k	274k	274k	524k	523k	62k	62k	125k	35k	274k	57k	61k	194k
BOOT	29k	59k	59k	256k	254k	28k	28k	35k	21k	58k	7k	29k	72k
GMM	33k	95k	101k	271k	236k	31k	30k	58k	23k	113k	14k	31k	86k
LDA	32k	89k	95k	238k	233k	33k	33k	56k	21k	96k	14k	30k	81k

Table 3: **Training Corpus Sizes** (number of selected instances) for zero-shot parsing experiments from target/proxy in-language treebanks (TARGET; where available), random sentence selection (RAND) and closest sentence selection (SENT), treebanks containing target genre (META), instances classified as target genre (BOOT), closest cluster selection (GMM and LDA).

SETUP	SWL ●	SA ■	KPV ■	TA ■	GL ■	YUE ●	CKT ●	FO w	TE ✗	MYV ■	QHE ✎	QTD ●	AVG
TARGET	40.66	38.74	26.70	75.83	85.51	—	—	58.78	91.26	—	73.62	66.75	61.98
RAND	22.81	47.06	25.97	72.14	84.68	49.70	29.39	71.66	83.73	36.88	40.63	58.97	51.97
SENT	24.47	44.98	31.69	71.28	84.63	51.11	31.95	71.92	83.03	41.73	40.19	58.85	52.99
META	24.94	44.62	25.77	72.26	84.26	47.91	22.66	70.54	82.06	36.67	19.83	57.93	49.12
BOOT	24.83	42.00	39.40	73.38	84.19	60.72	35.42	75.21	84.05	39.03	27.59	57.15	53.58
GMM	25.18	44.19	37.77	74.33	84.55	60.61	37.53	77.00	82.89	38.09	26.65	59.52	54.02
LDA	27.42	44.84	40.33	72.93	84.27	60.06	35.68	77.23	84.70	38.78	27.61	58.46	54.36

Table 4: **Unlabeled Attachment Scores** for zero-shot parsing experiments on test splits of target treebanks using training data from from target/proxy in-language treebanks (TARGET; where available), random sentence selection (RAND) and closest sentence selection (SENT), treebanks containing target genre (META), instances classified as target genre (BOOT), closest cluster selection (GMM and LDA).

SETUP	SWL ●	SA ■	KPV ■	TA ■	GL ■	YUE ●	CKT ●	FO w	TE ✗	MYV ■	QHE ✎	QTD ●	AVG
TARGET	0.71	0.54	0.77	1.16	0.24	—	—	1.32	0.97	—	0.26	1.10	0.79
RAND	1.60	0.46	0.16	0.72	0.09	1.33	0.89	1.02	0.64	1.09	0.55	0.55	0.76
SENT	2.13	2.00	0.58	1.76	0.18	0.67	0.27	0.63	0.92	0.37	0.37	0.91	0.90
META	0.90	0.75	0.73	1.24	0.27	0.41	1.19	0.82	0.42	0.44	0.44	0.73	0.73
BOOT	0.54	0.85	0.55	1.07	0.27	0.14	0.51	0.92	0.42	0.28	1.08	0.43	0.59
GMM	1.14	1.02	0.75	1.00	0.18	0.28	0.80	1.30	1.35	1.28	0.63	0.47	0.85
LDA	0.74	2.29	0.23	1.96	0.14	0.65	1.32	0.41	0.44	0.81	1.23	0.25	0.87

Table 5: **Standard Deviations of LAS** for zero-shot parsing experiments on test splits of target treebanks using training data from from target/proxy in-language treebanks (TARGET; where available), random sentence selection (RAND) and closest sentence selection (SENT), treebanks containing target genre (META), instances classified as target genre (BOOT), closest cluster selection (GMM and LDA).




Figure 4: **Selection Proportions** per target treebank and data selection method across all of UD. Zero instances were selected from shaded regions.