

Can Character-based Language Models Improve Downstream Task Performance in Low-Resource and Noisy Language Scenarios?

Arij Riabi^{1,2} Benoît Sagot¹ Djamé Seddah¹

¹ Inria Paris

² Sorbonne Université

{arij.riabi, benoit.sagot, djame.seddah}@inria.fr

Abstract

Recent impressive improvements in NLP, largely based on the success of contextual neural language models, have been mostly demonstrated on at most a couple dozen high-resource languages. Building language models and, more generally, NLP systems for non-standardized and low-resource languages remains a challenging task. In this work, we focus on North-African colloquial dialectal Arabic written using an extension of the Latin script, called *NArabizi*, found mostly on social media and messaging communication. In this low-resource scenario with data displaying a high level of variability, we compare the downstream performance of a character-based language model on part-of-speech tagging and dependency parsing to that of monolingual and multilingual models. We show that a character-based model trained on only 99k sentences of *NArabizi* and fine-tuned on a small treebank of this language leads to performance close to those obtained with the same architecture pre-trained on large multilingual and monolingual models. Confirming these results on a much larger data set of noisy French user-generated content, we argue that such character-based language models can be an asset for NLP in low-resource and high language variability settings.

1 Introduction

Current state-of-the-art monolingual and multilingual language models require large amounts of data to be trained, showing limited performance on low-resource languages (Howard and Ruder, 2018; Devlin et al., 2019). They lead to state-of-the-art results on most NLP tasks (Devlin et al., 2018; Raffel et al., 2020). In order to achieve high performance, these models rely on transfer learning architectures: the language models need to be trained on

large amounts of data (*pre-training*) to be able to transfer the acquired knowledge to a downstream task via *fine-tuning* on a relatively small number of examples, resulting in a significant performance improvement with respect to previous approaches. This dependency on large data sets for pre-training is a severe issue for low-resource languages, despite the emergence of large and successful multilingual pre-trained language models (Muller et al., 2021b). This is especially the case for languages with unusual morphological and structural features, which struggle to take advantage from similarities with high-resource, well represented languages such as Romance and Germanic languages.

In this work, we focus on one of such highly challenging languages, namely North-African dialectal Arabic. Its Latin transcription (*Arabizi*) displays a high level of linguistic variability¹, on top of scarce and noisy resource availability, making it a particularly challenging language for most NLP systems relying on pre-trained multilingual models (Muller et al., 2020).² To tackle the resource scarcity issue regarding Arabic dialects, Antoun et al. (2020) use BERT architecture (Devlin et al., 2019) to train a model on Arabic text to compare this approach to standard multilingual models. Indeed, Martin et al. (2020) show that fine-tuning a monolingual model leads to better results than fine-tuning a multilingual one, meaning that when fine-tuning is used there is no significant performance improvement from cross-lingual transfer during pre-training. However, such model is still pre-trained on sentences written in a single language and was not trained to handle the presence

¹Language variability, or language variation, is a term coming from socio-linguistics where, as stated by Nordquist (2019), it refers to regional, social or contextual differences in the ways that a particular language is used. These variations in user-generated content can be characterized through their prevalent idiosyncrasies when compared to canonical texts (Seddah et al., 2012a; Sanguinetti et al., 2020).

²Following Seddah et al. (2020), we refer to the Arabizi version of North-African Arabic dialects as *NArabizi*.

First version submitted on August 27th, 2021. Final on October 1st, 2021.

of multiple languages in the same sentence (code-switching), a frequent phenomenon in *NArabizi*.

However both monolingual and multilingual model approaches bear the risk of being limited by a subword tokenization-based vocabulary when facing out-of-domain training data language, especially in high-variability noisy scenarios (El Boukkouri et al., 2020; Clark et al., 2021), even though Muller et al. (2020) demonstrated a positive effect for *NArabizi* when using target language data to fine-tune a multilingual language model on its own objective function before pre-training.

Following a different approach, we investigate the use of a recently issued character-based language model (El Boukkouri et al., 2020) that was shown to display a remarkable robustness to lexical variation and noise when facing a new distant domain, namely biomedical. The pipeline we developed is simple and consists in fine-tuning this character-based model for several tasks in a noisy low-resource language scenario. We show that a character-based model trained on only 99k sentences of *NArabizi* and fined-tuned on a small treebank of the language leads to performance close to that obtained with the same architecture pre-trained on large multilingual and monolingual models (mBERT and CamemBERT).

Interestingly, we generalize this observation by using the same architecture on a much larger French user-generated Content treebank that exhibits similar language variability issues than *NArabizi*. In fact, pre-training a character-based model on 1% of the large-scale French instance of the multilingual corpus OSCAR leads to similar performance as a subword based model trained on the full corpus, showing that such character-based language model can reach similar performance levels and that the resulting models exhibit the same tolerance to noise as their much larger BERT counterparts. This demonstrates the value of such models in very scarce resource scenario. Our code and models are freely available.³

2 *NArabizi*: A Challenging Use Case for NLP in Low-resource Scenarios

As the official language of 25 countries, Arabic showcases a linguistic phenomenon called *diglossia* (Habash, 2010). It means that the speakers use

³<https://gitlab.inria.fr/ariabi/character-bert-ugc>

Modern Standard Arabic (MSA) for formal and official situations but use other forms of Arabic in informal situations. These dialectal forms constitute a dialect continuum with large variability from one country to the other. Arabic dialects are defined by their spoken form and often exhibit a lack of standardized spelling when written. When Arabic speakers produce written text in such dialects, they merely transcribe their spoken, colloquial language, which leads to different forms for the same word. Many users use the Latin script to express themselves online in their dialect (Seddah et al., 2020). In particular, they transcribe phonemes that cannot be straightforwardly mapped to a Latin letter using digits and symbols,⁴ with a high degree of variability at all levels; this is called *Arabizi*, with its North-African version called *NArabizi* by Seddah et al. (2020). For cultural and historical reasons, *NArabizi* also exhibits a high degree of code-switching with French and Amazigh (Amazouz et al., 2017), i.e. alternations between two or more languages during the conversation. Besides, the only available textual resources for Arabizi data are user-generated content, which is by itself inherently noisy (Foster, 2010; Seddah et al., 2012a; Eisenstein, 2013), making the production of supervised models, assuming the availability of labeled data, or even the collection of large pre-training data set a rather difficult task.

This data scarcity problem is often solved in the NLP literature using transfer learning: transferring knowledge learnt by large scale language models pre-trained on larger corpora. However, the use of Latin script makes it harder to transfer knowledge from language models trained on Arabic script corpora. For subword-based tokenization models, words are represented by a combination of subwords from a predefined list. When applied to a highly variable language with code-switching, a large vocabulary would be necessary to get a good covering of all possible orthographic variations which makes this approach less practical. Table 1 presents several examples of lexical variation within *NArabizi*. Interestingly, this variability also affects the code-switched vocabulary, which is mostly French in this case.

⁴For example, the digit “3” is often used to denote the ayin consonant, because it is graphically similar to its rendition in Arabic script.

GLOSS	ATTESTED FORMS	LANG
why	wa3lach w3alh 3alach	<i>NArabizi</i>
all	ekl kal kolach koulli kol	<i>NArabizi</i>
many	beaucoup boucoup bcp	Code-switched Fr.

Table 1: Examples of lexical variation in *NArabizi*. (From *Seddah et al., 2020*)

3 Related work

3.1 Transfer Learning for Low-resource Languages

Low-resource languages, by definition, face a lack of textual resources – annotated or not –, which makes it difficult for the NLP community to develop models and systems adapted to them (*Joshi et al., 2020*). The majority of the almost-7000 languages worldwide actually fall into the “low-resource” category. This makes the development of systems for low-resource languages necessary to widen the accessibility of NLP technology.

For deep learning approaches, which depend on the availability of large data sets, the solution to the low-resource problem comes from the idea of transfer learning. Early instances of cross-lingual transfer learning rely on non-contextualised word embeddings (*Ammar et al., 2016*). More recently, multilingual pre-trained language models (*Conneau et al., 2020*) has spread far and wide in NLP, enabling high-performance zero-shot cross-lingual transfer for numerous tasks and languages. The main idea is to exploit a large amount of unlabeled data to pre-train a model using a self-supervised task, such as masked language modeling (*Lample and Conneau, 2019; Vania et al., 2019*). This pre-trained model is then fine-tuned on a much smaller annotated data set and used for another language, domain or task. Strategic knowledge sharing has been shown to improve the performance on downstream tasks and languages (*Gururangan et al., 2020*). Therefore, this technique is crucial for multilingual applications, as most of the world’s languages lack large amount of labeled data (*Conneau et al., 2019; Eisenschlos et al., 2019; Joshi et al., 2020*). However the performance of multilingual language model on low-resource languages is still limited compared to other languages since they are naturally under-sampled during the training process (*Wu and Dredze, 2020*).

To improve performance on a specific low-resource languages, there are two possibilities. Either to attempt to train a language model on it from

scratch despite the scarceness of data; or fine-tune a pre-trained multilingual language model on the low-resource language corpus, also called *cross-lingual transfer learning* (*Muller et al., 2021a*). The first option can sometimes lead to decent performance, provided that the training corpus is diverse enough (*Martin et al., 2020*).

When following the fine-tuning approach, unsupervised methods can be implemented to facilitate the transfer of knowledge (*Pfeiffer et al., 2020*). The most widely used unsupervised fine-tuning task is masked language modeling (MLM). This system has proven its efficiency between languages that have already been seen in the training corpus (*Pires et al., 2019*); it is still a challenge for unseen languages, especially low-resource ones. However, *Muller et al. (2020)* achieved promising results by performing unsupervised fine-tuning on small amounts of *NArabizi* data. We follow their approach by comparing the performance of our pipeline in two setups: MODEL+MLM+TASK and MODEL+TASK. We describe these setups in more details in section 6.

3.2 Tokenization & Character-based models

Standard language models rely on a subword-based approach to process tokens in a sequence (*Kudo, 2018a*). This allows the model to handle any word unseen in the training data, working in an “open-vocabulary setting,” where words are represented by a combination of subwords from a pre-defined list. On top of alleviating the issue of out-of-vocabulary words, this approach allows the model to handle sequences written in a language unseen during training, as long as it uses the same script. Therefore, subword tokenization is a crucial feature of state-of-the-art models in NLP. But its suitability for all types of data has always been questioned. While splitting texts into subwords based on their frequencies works well for English, models using this kind of tokenization struggle with noise, whether it is naturally present in the data (*Sun et al., 2020*) or artificially generated to challenge the model (*Pruthi et al., 2019*). Moreover, language models that use subword-based tokenization struggle to represent rare words (*Schick and Schütze, 2020*). Many research projects have focused on improving subword tokenization. For example, *Wang et al. (2021)* suggested a multi-view subword regularization based on the sampling of multiple segmentations of the input text, based on

the work of Kudo (2018b).

Other parallel efforts bid on character-based models. For example, El Boukkouri et al. (2020) proposed a possible solution to get a better tokenization system more resilient to orthographic variations and noise in the data set by using a character-level model, inspired by a previous word-level open-vocabulary system (Peters et al., 2018a). This new model gets better results than vanilla BERT on multiple tasks from the medical domain. Furthermore, the authors claim that it is more robust to noise and misspellings. In the same vein, Ma et al. (2020a) combined character-aware and subword-based information to improve robustness to spelling errors. This initiated a new wave of tokenizer-free models based on characters or bytes (Tay et al., 2021; Xue et al., 2021; Clark et al., 2021).

The question of knowing if character-based language models can handle high language variability since they are supposed to be resilient to noise and spelling variations is crucial when dealing with non-normalized dialects and non-canonical forms of language as found on many user-generated content platforms. This is why we focus on this work on the analysis of the performance of character-based models on several user-generated content data sets that we now describe.

4 Data sets

In this section, we describe the data sets we use to evaluate our pipeline on our downstream tasks, namely Part-Of-Speech (POS) tagging and dependency parsing.

***NArabizi* Data Set** We use the *NArabizi* treebank (Seddah et al., 2020), containing about 1500 sentences randomly sampled from the romanized Algerian dialectal Arabic corpus of Cotterell et al. (2014) and from a small corpus of lyrics coming from Algerian dialectal Arabic songs popular among the younger generation. This treebank is manually annotated with morpho-syntactic information (parts-of-speech and morphological features), together with glosses and code-switching labels at the word level, as well as sentence-level translations. Moreover, this treebank also contains 36% of French tokens. Within the *NArabizi* annotated corpus,⁵ In addition to labeled data, The *NArabizi* treebanks provides about 2 millions words, 99k

⁵<http://almanach-treebanks.fr/NArabizi>

Treebank	# Tokens	# Sentences	Genres
GSD	389,363	16,342	Blogs, News Reviews, Wiki
Sequoia	68,615	3,099	Medical, News Non-fiction, Wiki
Spoken	34,972	2,786	Spoken
FSMB	56,009	4,055	Twitter, Facebook Web Forums

Table 2: Statistics on the treebanks used in our POS tagging and dependency parsing experiments.

sentences, of unlabeled data collected from various sources.⁶ We use this corpus for unsupervised fine-tuning.

French Data Sets We use the following Universal Dependencies, UD, (Nivre et al., 2020) version of the following treebanks: French GSD (McDonald et al., 2013), Sequoia (Candito and Seddah, 2012) and Spoken, an automatic conversion of the Rhapsodie corpus (Lacheret et al., 2014) to the UD annotation scheme.

For our experiments on noisy UGC treebank, we use an extension of the French Social Media Bank, FSMB⁷ (Seddah et al., 2012a): a treebank of French sentences coming from various social media sources only available either in constituent trees or in the native French Treebank dependency annotation scheme (Candito et al., 2010), along with the Sequoia original treebank that we use with the same annotation scheme for compatibility in our French UGC experiments. A brief overview of the size and content of each treebank can be found in Table 2.

Pre-training Data Sets Note that in some of our experiments, we use a fraction, 1% of the deduplicated French instance of the Oscar corpora (Suárez et al., 2019), about 380M words, as a source of unlabeled data to be either mixed with *NArabizi* pre-training data (as an 2.5M words additional sample) or used as pre-training data for characterBert when evaluated on French UGC (whole 1%). Statistics on those data sets are presented on Table 3.

5 Model

CharacterBERT (El Boukkouri et al., 2020) is a character-based variant of BERT that replaces the WordPiece embedding matrix with multiple CNN

⁶The original data set provides 50k sentences of clean *NArabizi* sentences and an additional 49k sentences of more noisy data, we use here a concatenation of both.

⁷We use a shuffled version of the treebank split into a train set of about 2 000 sentences and a dev and test set of about 1 000 sentences each.

	# Tokens	# Sentences	Language
99k Narabizi	2.527k	99k	ar-dz
1% Oscar	318.715k	9.342k	fr
10% Oscar	1.885.351k	55.261k	fr
100% Oscar	23.209.872k	558.092k	fr

Table 3: pre-training data set statistics.

and a highway layer (Srivastava et al., 2015). This method for encoding token representations is inspired from ELMo (Peters et al., 2018b), one of the first pre-trained language models for transfer learning. It generates a context-independent representation from character embeddings and feeds them to transformer encoder layers, similarly to BERT architecture. Therefore, it produces a single embedding for any input token and does not need a WordPiece vocabulary. This avoids having an inconstant number of subword vectors for each word. We choose this model since its robustness to noise, when tested on biomedical domain by El Boukkouri et al. (2020), can lead to interesting result when facing our noisy experiment data. It is the first character-based and BERT-like model, along with CharBERT (Ma et al., 2020b), that uses both character and subword embeddings and has the advantage of being publicly available.⁸ Note that we retrain CharacterBert from scratch on our data sets and do not make use of any of its available pretrained at models any point in our experiments.

6 Experiments

In this section, we present our fully-supervised and semi-supervised baselines. We also evaluate different fine-tuning strategies combined with two layers configurations. We use various embedding models that we contrast with the CharacterBERT model.

Baseline Models For our fully-supervised baseline, we use FastText embeddings (Joulin et al., 2016) trained from scratch on our treebank training sets and used as input for our downstream tasks without any special treatment.

To measure the effectiveness of using a contextualized character-based language model, we compare its performance to subword based language models, both monolingual and multilingual, that constitute the basis of our semi-supervised baselines. For multilingual subword based language

⁸<https://github.com/helboukkouri/character-bert-pre-training>

	UPOS	UAS	LAS
No external embeddings	57.61	55.48	39.32

Table 4: Pos-Tagging and Parsing Baseline Results for *NArabizi* test set.

model, we use mBERT, the multilingual version of BERT (Devlin et al., 2018). It is trained for 104 different languages on Wikipedia data, including French and Arabic, languages for which Muller et al. (2020) showed that, to a certain extent, they could transfer to *NArabizi*. For our monolingual model we use CamemBERT (Martin et al., 2020) which is a contextualized language model based on the RoBERTa model (Liu et al., 2019) trained and optimized specifically for French.

MODEL+TASK We use the implementation of the Biaffine graph parser (Dozat and Manning, 2016) from Grobol and Crabbé (2021); they adapted it to handle several input sources, such as BERT representations. The parser performs its own tagging using a multi-layer perceptron. The word representations are a concatenation of word embeddings and tag embeddings learned together with the model parameters on the treebank training data. We fine-tune the overall model by back-propagating through the average of all sub-tokens of a word. Baseline results, without any external embeddings, for this parser are provided in Table 4.

MODEL+MLM+TASK Before fine-tuning the model on the downstream task, we perform language adaptation by fine-tuning it in a self-supervised fashion with the MLM loss. We use the *NArabizi* raw data; we train the model for 20 epochs, keeping the best model obtained at the end. The evaluation is done using the MLM log likelihood, with 10% of the data kept for validation.

Fine-Tuning Strategy In addition to using only the *last layer* (cf. Appendix A where we conducted several experiments to explore the effect of different layer configurations in our downstream tasks), we also test two options for the layers aggregation. The first one is simply taking the *mean* of selected layers. The second is *scalar mix*, introduced by ELMo (Peters et al., 2018b): a convex combination of the transformer layers where the weights are learnt.

For each of these configurations, we test two setup: *with* and *without* training the transformer

Fine-tuning Strategy	MODEL+TASK			MODEL+MLM+TASK			Fine-tuning Strategy	MODEL+TASK			MODEL+MLM+TASK		
	UPOS	UAS	LAS	UPOS	UAS	LAS		UPOS	UAS	LAS	UPOS	UAS	LAS
last-layer-ft	81.14	72.59	60.35	83.08	73.77	62.00	last-layer-ft	80.48	69.19	57.89	84.55	73.82	62.67
last-layer-fz	56.05	62.48	43.95	70.23	67.86	53.45	last-layer-fz	65.08	62.95	45.84	79.44	71.55	58.98
mean-ft	80.86	73.06	61.11	84.22	72.97	61.58	mean-ft	84.40	72.83	61.81	84.50	71.31	60.96
mean-fz	65.41	64.27	48.11	73.53	68.05	53.92	mean-fz	69.47	64.89	49.34	78.40	69.71	57.61
scalar-mix-ft	80.15	70.46	58.08	83.65	74.62	62.62	scalar-mix-ft	80.25	69.19	56.14	84.31	72.78	61.67
scalar-mix-fz	58.74	63.80	46.93	61.20	67.25	52.84	scalar-mix-fz	71.79	66.73	52.17	78.83	72.07	59.17

Fine-tuning Strategy	(a) CamemBERT			Sample Oscar			(b) mBERT		
	NArabizi			99k			NArabizi + Oscar		
	UPOS	UAS	LAS	UPOS	UAS	LAS	UPOS	UAS	LAS
last-layer-ft	81.19	70.56	58.65	78.83	69.52	56.33	80.67	69.90	57.75
last-layer-fz	75.71	66.87	53.69	71.74	65.31	51.75	75.38	68.57	55.15
mean-ft	81.00	70.70	58.84	78.78	69.52	55.81	80.91	69.57	57.18
mean-fz	77.08	68.05	55.53	72.35	66.54	51.98	77.22	69.33	56.14
scalar-mix-ft	79.96	69.47	57.84	80.62	68.76	56.76	80.15	69.90	57.56
scalar-mix-fz	75.90	68.10	54.58	71.69	66.16	51.80	77.60	69.33	55.91

(c) CharacterBERT (MODEL+TASK)

Table 5: Performances of the models on the *NArabizi* treebank using **different fine-tuning strategy** (We use *ft* to indicate that the embeddings are fine-tuned for the tasks, while *fz* is used when the embeddings are frozen during the fine-tuning step).

model weights during the fine-tuning. We call the first strategy *frozen representation*, noted with a *-fz* suffix in our results, where the language model is used to extract meaningful features that consist of contextual embeddings. In the second strategy, the contextual embedding extractor – that is, the pre-trained language model – is fine-tuned on the downstream task alongside the task-specific component, noted with a *-ft* suffix in our results.

7 Results and Discussion

In this section, we compare all the systems presented before on the *NArabizi* and French treebanks. We especially focus on the impact of corpus size on the models’ performance.

7.1 Experiments on *NArabizi* treebank

Scores are reported as triplets describing UPOS/UAS/LAS results. Highlighted cells marks the best results column-wise while bold marks best results row-wise

We report in Table 5 the scores for the different fine-tuning strategies for the three models CamemBERT, mBERT and CharacterBERT.

Additional insights on the extraction of a representation from the different layers are provided in the Appendix, Table 9, with analysis on the effect of the combination of different subsets of layers on the accuracy.

CharacterBERT_{NArabizi} performs better overall in the Model+TASK setting Looking at table

5, we notice that CharacterBERT_{NArabizi} significantly outperforms mBERT and CamemBERT in the MODEL+TASK setup without MLM for almost all the configurations, except for the two configurations mean-ft (when using the mean of all the fine-tuned layers) and scalar-mix-ft (when using the scalar-mix of all the layers) where mBERT and CamemBERT show slightly better performance than CharacterBERT_{NArabizi}. In this latter case, mBERT outperforms CharacterBERT_{NArabizi} on the *NArabizi* data set with accuracy differences of (3.40/2.13/2.96) for (UPOS/UAS/LAS) using mean-ft, which is only significant for UAS with p-value<0.05.⁹ The same observation can be done for CamemBERT in the MODEL+TASK setting which outperforms CharacterBERT_{NArabizi} with accuracy differences of (-0.14/2.36/2.26).

The same is observed for the scalar-mix-ft option but the differences are not significant with p-value<0.05 for UAS and LAS. Besides these two configurations, CharacterBERT_{NArabizi} always outperforms the other two models without MLM. The most notable difference using CharacterBERT_{NArabizi} is when the latter outperforms CamemBERT with accuracy differences of (19.66/4.39/9.74) and mBERT with (10.63/3.92/7.854) in the last-layer-fz setting with a p-value<0.05 for both comparisons.

⁹We tested the statistical significance using the publicly available Dan Bikel’s code at <https://github.com/tdozat/Parser-v1>

Adding MLM reverses the trend However, when we compare CharacterBERT_{NArabizi} to CamemBERT+MLM and mBERT+MLM, we see that while they both generally outperform CharacterBERT_{NArabizi}, there are some settings where CharacterBERT_{NArabizi} still gets better scores. mBERT+MLM gets the best scores for all of the configurations we test among the three models. While both CamemBERT and mBERT perform comparably, mBERT outperforms CharacterBERT_{NArabizi} when it is used with MLM. Some of the best performance are achieved in the last-layer-ft setting with scores of (83.08/73.77/62.00) for CamemBERT and (84.55/73.82/62.67) for mBERT and lower scores of (81.19/70.56/58.6) for CharacterBERT_{NArabizi} (both models outperform CharacterBERT_{NArabizi} in this case).

In other settings, CharacterBERT_{NArabizi} is still competitive with the two other MLM-pretrained models, as illustrated by the last-layer-fz setting with scores of (70.23/67.86/53.45) for CamemBERT, and (75.71/66.87/53.69) for CharacterBERT_{NArabizi} and (79.44/71.55/58.98) for mBERT. **Still, the general tendency in the Model+MLM+TASK setting is that mBERT outperforms CharacterBERT_{NArabizi} when used with MLM and CamemBERT exhibits performance similar to CharacterBERT_{NArabizi} in the same setting.** This is in contrast with the earlier comparison without MLM pre-training where both CamemBERT and mBERT reached scores lower than those of CharacterBERT_{NArabizi}.

This observation confirms the findings of [Muller et al. \(2021a\)](#) regarding the positive impact of unsupervised fine-tuning for BERT models even if the language is not one of the pre-training languages.

CharacterBERT pretrained on 1% of Oscar performs roughly the same than CamemBERT+Task If we compare the performance of CamemBERT in the MODEL+TASK setting to CharacterBERT trained on sub-sample of Oscar, we see that both models are comparable. Typically, CamemBERT seems to outperform CharacterBERT in some settings like last-layer-ft where the latter records (78.83/69.52/56.33) for UPOS, UAS and LAS scores respectively while CamemBERT records higher at (81.14/72.59/60.35). In other settings, CharacterBERT seems to outperform CamemBERT. In the mean-fz setting for instance, CharacterBERT has scores of (72.35/66.54/51.98)

which surpasses the (69.47/64.89/49.34) scores of CamemBERT. No clear conclusion can be drawn about the best use of one model over the other in the different settings since they all display competitive scores. This is essentially due to the fact that CamemBERT is trained on the full Oscar data set, while CharacterBERT is trained on just 0.01% of it. In addition, the test set is made of only 140 *NArabizi* sentences, making any interpretation of the results difficult. These two reasons make difficult the drawing of concrete conclusions on the performance of both models compared to each other. Therefore in the next section, we will evaluate the models using the best fine-tuning strategy on French treebanks with a larger evaluation set and a CharacterBERT trained on 1% Oscar.

7.2 The Impact of Data Size: Experiments on French treebanks

In-domain experiments For our experiments on French treebanks, we keep the best setup from our previous results and use the last layer of the model fine-tuned for the task.

In table 6, we report the scores obtained by both models CamemBERT and CharacterBERT on three French treebanks, compared to our parser without external embeddings from pre-trained models, which we consider as our baseline for those experiments. All systems using CamemBERT or CharacterBERT models outperform our baseline. Both models have competitive scores. For example, when tested on the GSD treebank, our baseline obtains scores of (96.35/89.98/86.96) while CamemBERT and CharacterBERT obtain both better, yet relatively close scores, with (98.53/95.74/94.17) for the CamemBERT-based model and (98.07/95.22/93.50). The same observations apply for the other two treebanks SEQUOIA and SPOKEN. These results are coherent with the comparisons done earlier for *NArabizi* where both CharacterBERT and CamemBERT models produced comparable scores even though CharacterBERT was trained on only 0.01% of the full data set on which CamemBERT was trained on. Our results on those much larger data sets corroborate then our *NArabizi* models results and confirm the interest of using characterBert-based models in our scenarios.

Extremely noisy user-generated content experiments In table 7, we compare both models CamemBERT and CharacterBERT in two different

Model-LayerConfig	GSD			SEQUOIA			SPOKEN		
	UPOS	UAS	LAS	UPOS	UAS	LAS	UPOS	UAS	LAS
No pre-training	96.35	89.98	86.96	92.57	82.28	77.72	83.60	67.65	58.04
CamemBERT (100% Oscar)	98.53	95.74	94.17	99.23	95.78	94.59	97.41	88.15	83.00
CharacterBERT (1% Oscar)	98.07	95.22	93.50	99.27	94.98	93.80	96.62	86.82	81.22

Table 6: POS and dependency parsing scores on 3 French treebanks.

Model-LayerConfig	Dev			Test			
	UPOS	UAS	LAS	UPOS	UAS	LAS	
FSMB train	No pre-training	81.24	70.04	59.46	81.62	69.19	59.17
	CamemBERT	95.34	87.01	81.56	95.48	87.47	82.66
	CharacterBERT	95.08	85.99	80.51	95.19	86.26	81.26
Sequoia train	No pre-training	71.50	59.04	47.30	72.79	59.92	48.81
	CamemBERT	89.47	81.80	74.33	90.10	82.68	75.85
	CharacterBERT	90.12	81.79	74.43	90.68	82.39	75.39

Table 7: POS and dependency parsing scores on the FSMB.

settings. In the first one, both models are fine-tuned on the FSMB training set. In the second one, the models are trained on the Sequoia training set. The performance of these models is to be compared – as in table 7 – to the performance of a parser without external embeddings from a pre-trained model. As expected, both models outperform the baseline and – similarly to previous results – are competitive in their respective obtained scores.

In the FSMB training setting, CamemBERT achieves scores of (95.34/87.01/81.56) when tested on the development part of the data set, and CharacterBERT achieves close scores of (95.08/85.99/80.51) on the same data set. Both scores outperform the ones obtained by the baseline of (81.24/70.04/59.46) also on the development set. The same behavior can be observed in the Sequoia training setting where CamemBERT gets scores of (89.47/81.80/74.33) and CharacterBERT scores of (90.12/81.79/74.43) on the development set: both are still higher than the baseline scores. This comparison is still valid when we consider the test set results. For reasons probably tied to the random sampling done when splitting the data set, the score ranking are reversed in the test set, yet results are very similar and the slight differences between the two models results are not statistically significant. This shows the actual effectiveness of having a character-based model trained on only a small fraction of its “classic” BERT counterpart when facing noisy user-generated content from a data set that was proven to be much more noisy than many other similar data sets (Rosales Núñez et al., 2019).

	UPOS	UAS	LAS	%Oscar
	<i>FSMB fine-tuned (in-domain)</i>			
CamemBERT	95.48	87.47	82.66	100
CamemBERT _{4gb}	95.13	85.73	80.72	2.38
Character-BERT	95.19	86.26	81.26	1
	<i>Sequoia fine-tuned (out domain)</i>			
CamemBERT	90.10	82.68	75.85	100
CamemBERT _{4gb}	90.69	82.29	75.83	2.38
Character-BERT	90.68	82.39	75.39	1

Table 8: CharacterBert model performance compared with a small CamemBERT (4gb) model on the FSMB test set in in-domain and out-of-domain fine-tuning scenarios. *Full-size Camembert results are reported here for reference.*

8 Discussion

In this work, we evaluate the benefits of using a character-based model in low-resource scenarios. Our results show that training such a model from scratch on much fewer data gives similar performance to a multilingual BERT adapted to the language using the same amount of data.

Overall, our observations confirm the findings of El Boukkouri et al. (2020) regarding the robustness to noise and misspellings of the CharacterBERT model. We showed that the model has competitive performance on noisy French UGC data when trained on only a fraction of the OSCAR corpus compared to CamemBERT trained on the full corpus and when trained on corpora containing about 1M words in the extremely noisy and low-resource case of NArabizi. This is consistent with the findings of Martin et al. (2020) and Micheli et al. (2020), who showed that MLM could already learn a lot from pre-training on smaller data set. Extending this investigation by training on a larger

amount of data could help to explore the ability of the model to handle highly variable noisy data.

However, one could question the usefulness of such Character-BERT based models if small Bert-based models were available on the same domain. To build an answer to that question, we conducted a quick set of experiments comparing our characterBert model trained on 1% of Oscar with the off-the-shelf Camembert version trained on 4gb of the Oscar corpus French instance (2.38% of the full corpus) and which was shown to perform almost as well as the full model (Martin et al., 2020) on many downstream tasks. Both models were fine-tuned according to our MODEL+Task architecture on either the FSMB or the Sequoia treebank, allowing us to evaluate their in-domain and out-of-domain performance. Results on Table 8 confirm the effectiveness of our characterBert model with overall better results than CamemBERT_{4gb} in the in-domain scenario and similar, if not slightly better in the out-of-domain scenario, except for the labeled attachment score (75.83 vs 75.39). The fact that CamemBERT_{4gb} was trained on more than twice as much data and with 200k pre-training steps while the characterBert pre-training stopped below 20k steps probably explains this small discrepancy but further investigations are needed with a fully parallel setting where both characterBert and CamemBERT are pretrained on the same amount of data and the same hyper-parameters. The take-home message from this in-domain experiments is that CharacterBert seems to be able to better capture at least some of the UGC idiosyncrasies that are prevalent in the FSMB (Seddah et al., 2012b) than its Bert-based counterparts. This was also shown by Rosales Núñez et al. (2021) in the context of character-based neural machine translation. Interestingly, their results showed that transformer-based models with subword tokenization also exhibit strong robustness to a certain type of lexical noise. This behavior has been very recently demonstrated by Itzhak and Levy (2021) and could explain why the BERT-based models we tested performed so well in our experiments. The key seems to be relying on the ability of the subword distribution to model some forms of lexical variations. Much more experiments are needed to clearly investigate in what circumstances, besides noisy and resource-scarce scenarios, characterBERT models bring in a decisive advantage.

Our results are based on the evaluation of two

low-level tasks. Therefore, it would be interesting to see if they can be generalized to other – e.g. more semantic – tasks, as additional experiments on model layers configuration showed that most of the important information is captured early in the layers of the model (cf. Appendix A).

Regarding the specific case of Arabic dialects written in Arabizi, a recent BERT-based model have been pretrained on 7 millions Egyptian tweets and displayed effective results on a sentiment analysis task (Baert et al., 2020). Another very recent model, at the date of writing, was pre-trained on 4 millions Algerian tweets and also demonstrated interesting results on sentiment analysis (Abdaoui et al., 2021). Unfortunately, the authors did not perform any experiments on the Narabizi data set, making thus the comparison with our work not straightforward. It would be of course interesting to evaluate the interoperability between these new data sets and the *NArabizi* resources we used to produce our models. Head to head comparisons between these models and ours could be of value of course but we believe that given the shortcomings of finding enough data to pretrain large models for dialects, it would be probably better to first consolidate a large enough common pre-training data set and then work on model performance. We leave this for future work.

We showed that CharacterBert models trained on very little data could provide an interesting alternative to large multilingual and monolingual models in resource-scarce and noisy scenarios. This is why we release all the code, data and models to reproduce our experiments, hoping our work will favor the rise of efficient robust NLP models for under-resourced languages, domains and dialects.¹⁰

Acknowledgments

We thank the reviewers for their very valuable feedback. The first author was partly funded by Benoît Sagot’s chair in the PRAIRIE institute funded by the French national agency ANR as part of the “Investissements d’avenir” programme under the reference ANR-19-P3IA-0001. This work also received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 101021607 and from the French Research Agency via the ANR ParSiTi project (ANR-16-CE33-0021).

¹⁰<https://gitlab.inria.fr/ariabi/character-bert-ugc>

References

- Amine Abdaoui, Mohamed Berrimi, Mourad Oussalah, and Abdelouahab Moussaoui. 2021. [Dziribert: a pre-trained language model for the algerian dialect](#).
- Djegdjiga Amazouz, Martine Adda-Decker, and Lori Lamel. 2017. Addressing code-switching in french/algerian arabic speech. In *Interspeech 2017*, pages 62–66.
- Waleed Ammar, George Mulcaire, Yulia Tsvetkov, Guillaume Lample, Chris Dyer, and Noah A. Smith. 2016. Massively multilingual word embeddings. *ArXiv*, abs/1602.01925.
- Wissam Antoun, Fady Baly, and Hazem Hajj. 2020. Arabert: Transformer-based model for arabic language understanding. *arXiv preprint arXiv:2003.00104*.
- Gaétan Baert, Souhir Gahbiche, Guillaume Gadek, and Alexandre Pauchet. 2020. [Arabizi language models for sentiment analysis](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 592–603, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Marie Candito, Benoît Crabbé, and Pascal Denis. 2010. [Statistical French dependency parsing: Treebank conversion and first results](#). In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta. European Language Resources Association (ELRA).
- Marie Candito and Djamé Seddah. 2012. Le corpus sequoia: annotation syntaxique et exploitation pour l’adaptation d’analyseur par pont lexical. In *TALN 2012-19e conférence sur le Traitement Automatique des Langues Naturelles*.
- Jonathan H. Clark, Dan Garrette, Iulia Turc, and John Wieting. 2021. [Canine: Pre-training an efficient tokenization-free encoder for language representation](#).
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- 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.
- Ryan Cotterell, Adithya Renduchintala, Naomi Saphra, and Chris Callison-Burch. 2014. An Algerian Arabic-French code-switched corpus. In *Workshop on Free/Open-Source Arabic Corpora and Corpora Processing Tools Workshop Programme*, page 34.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- 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. 2016. Deep biaffine attention for neural dependency parsing. *arXiv preprint arXiv:1611.01734*.
- Julian Martin Eisenschlos, Sebastian Ruder, Piotr Czapla, Marcin Kardas, Sylvain Gugger, and Jeremy Howard. 2019. Multifit: Efficient multilingual language model fine-tuning. *arXiv preprint arXiv:1909.04761*.
- Jacob Eisenstein. 2013. What to do about bad language on the internet. In *Proceedings of the 2013 conference of the North American Chapter of the association for computational linguistics: Human language technologies*, pages 359–369.
- Hicham El Boukkouri, Olivier Ferret, Thomas Lavergne, Hiroshi Noji, Pierre Zweigenbaum, and Jun’ichi Tsujii. 2020. [CharacterBERT: Reconciling ELMo and BERT for word-level open-vocabulary representations from characters](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6903–6915, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Jennifer Foster. 2010. [“cba to check the spelling”: Investigating parser performance on discussion forum posts](#). In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 381–384, Los Angeles, California. Association for Computational Linguistics.
- Loïc Grobol and Benoît Crabbé. 2021. [Analyse en dépendances du français avec des plongements contextualisés](#). In *Actes de la 28ème Conférence sur le Traitement Automatique des Langues Naturelles*.
- 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.
- Nizar Habash. 2010. *Introduction to Arabic Natural Language Processing*. Morgan and Claypool.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*.
- Itay Itzhak and Omer Levy. 2021. [Models in a spelling bee: Language models implicitly learn the character composition of tokens](#).
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. What does bert learn about the structure of language? In *ACL 2019-57th Annual Meeting of the Association for Computational Linguistics*.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The state and fate of linguistic diversity and inclusion in the NLP world](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, H erve J egou, and Tomas Mikolov. 2016. Fasttext. zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*.
- Taku Kudo. 2018a. [Subword regularization: Improving neural network translation models with multiple subword candidates](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 66–75, Melbourne, Australia. Association for Computational Linguistics.
- Taku Kudo. 2018b. [Subword regularization: Improving neural network translation models with multiple subword candidates](#).
- Anne Lacheret, Sylvain Kahane, Julie Beliao, Anne Dister, Kim Gerdes, Jean-Philippe Goldman, Nicolas Obin, Paola Pietrandrea, and Atanas Tchobanov. 2014. Rhapsodie: a prosodic-syntactic treebank for spoken french. In *Language Resources and Evaluation Conference*.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *arXiv preprint arXiv:1901.07291*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Wentao Ma, Yiming Cui, Chenglei Si, Ting Liu, Shijin Wang, and Guoping Hu. 2020a. [Charbert: Character-aware pre-trained language model](#). *Proceedings of the 28th International Conference on Computational Linguistics*.
- Wentao Ma, Yiming Cui, Chenglei Si, Ting Liu, Shijin Wang, and Guoping Hu. 2020b. [Charbert: Character-aware pre-trained language model](#). *arXiv preprint arXiv:2011.01513*.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Su arez, Yoann Dupont, Laurent Romary,  Eric de la Clergerie, Djam e Seddah, and Beno t Sagot. 2020. [CamemBERT: a tasty French language model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online. Association for Computational Linguistics.
- Ryan McDonald, Joakim Nivre, Yvonne Quirnbach-Brundage, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith Hall, Slav Petrov, Hao Zhang, Oscar T ackstr om, et al. 2013. Universal dependency annotation for multilingual parsing. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 92–97.
- Vincent Micheli, Martin d’Hoffschmidt, and Fran ois Fleuret. 2020. [On the importance of pre-training data volume for compact language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7853–7858, Online. Association for Computational Linguistics.
- Benjamin Muller, Antonios Anastasopoulos, Beno t Sagot, and Djam e Seddah. 2021a. [When being unseen from mBERT is just the beginning: Handling new languages with multilingual language models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 448–462, Online. Association for Computational Linguistics.
- Benjamin Muller, Yanai Elazar, Beno t Sagot, and Djam e Seddah. 2021b. [First align, then predict: Understanding the cross-lingual ability of multilingual BERT](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2214–2231, Online. Association for Computational Linguistics.
- Benjamin Muller, Beno t Sagot, and Djam e Seddah. 2020. Can multilingual language models transfer to an unseen dialect? a case study on north african arabizi. *arXiv preprint arXiv:2005.00318*.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Haji c, 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.
- Richard Nordquist. [Linguistic variation](#) [online]. 2019.

- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018a. [Deep contextualized word representations](#). 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 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018b. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020. [MAD-X: An Adapter-Based Framework for Multi-Task Cross-Lingual Transfer](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7654–7673, Online. Association for Computational Linguistics.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Danish Pruthi, Bhuwan Dhingra, and Zachary C. Lipton. 2019. [Combating adversarial misspellings with robust word recognition](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5582–5591, Florence, Italy. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67.
- José Carlos Rosales Núñez, Djamé Seddah, and Guillaume Wisniewski. 2019. [Comparison between NMT and PBSMT performance for translating noisy user-generated content](#). In *Proceedings of the 22nd Nordic Conference on Computational Linguistics*, pages 2–14, Turku, Finland. Linköping University Electronic Press.
- José Carlos Rosales Núñez, Guillaume Wisniewski, and Djamé Seddah. 2021. Noisy ugc translation at the character level: Revisiting open-vocabulary capabilities and robustness of char-based models. In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2022)*, Punta Cana, Dominican Republic.
- Manuela Sanguinetti, Lauren Cassidy, Cristina Bosco, Özlem Çetinoglu, Alessandra Teresa Cignarella, Teresa Lynn, Ines Rehbein, Josef Ruppenhofer, Djamé Seddah, and Amir Zeldes. 2020. [Treebanking user-generated content: a UD based overview of guidelines, corpora and unified recommendations](#). *CoRR*, abs/2011.02063.
- Timo Schick and Hinrich Schütze. 2020. Rare words: A major problem for contextualized embeddings and how to fix it by attentive mimicking. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8766–8774.
- Djamé Seddah, Farah Essaidi, Amal Fethi, Matthieu Futral, Benjamin Muller, Pedro Javier Ortiz Suárez, Benoît Sagot, and Abhishek Srivastava. 2020. [Building a user-generated content North-African Arabizi treebank: Tackling hell](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1139–1150, Online. Association for Computational Linguistics.
- Djamé Seddah, Benoit Sagot, Marie Candito, Virginie Mouilleron, and Vanessa Combet. 2012a. [The French Social Media Bank: a treebank of noisy user generated content](#). In *Proceedings of COLING 2012*, pages 2441–2458, Mumbai, India. The COLING 2012 Organizing Committee.
- Djamé Seddah, Benoit Sagot, Marie Candito, Virginie Mouilleron, and Vanessa Combet. 2012b. The french social media bank: a treebank of noisy user generated content. In *COLING 2012-24th International Conference on Computational Linguistics*.
- Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. 2015. Highway networks. *arXiv preprint arXiv:1505.00387*.
- Pedro Javier Ortiz Suárez, Benoît Sagot, and Laurent Romary. 2019. Asynchronous pipeline for processing huge corpora on medium to low resource infrastructures. In *7th Workshop on the Challenges in the Management of Large Corpora (CMLC-7)*. Leibniz-Institut für Deutsche Sprache.
- Lichao Sun, Kazuma Hashimoto, Wenpeng Yin, Akari Asai, Jia Li, Philip Yu, and Caiming Xiong. 2020. [Adv-bert: Bert is not robust on misspellings! generating nature adversarial samples on bert](#).
- Yi Tay, Vinh Q. Tran, Sebastian Ruder, Jai Gupta, Hyung Won Chung, Dara Bahri, Zhen Qin, Simon Baumgartner, Cong Yu, and Donald Metzler. 2021. [Charformer: Fast character transformers via gradient-based subword tokenization](#).
- 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. *arXiv preprint arXiv:1909.02857*.
- Xinyi Wang, Sebastian Ruder, and Graham Neubig. 2021. [Multi-view subword regularization](#).
- Shijie Wu and Mark Dredze. 2020. [Are all languages created equal in multilingual BERT?](#) In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 120–130, Online. Association for Computational Linguistics.

Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2021. [Byt5: Towards a token-free future with pre-trained byte-to-byte models.](#)

A Layer Configuration Experiments

We focus on the scalar mix strategy to study the effect of the combination of different subsets of layers on the accuracy of the downstream task, instead of only aggregating all layers. For example, it has been shown that higher layers contain more semantic information, while lower layers contain more syntactic information (Jawahar et al., 2019). We compare several layer configurations, that is different subsets of the transformer from which we get the sentence representation.

The effect of the Layer Configuration We report in tables 9a, 9b and 9c the scores for the different layers combinations for CamemBERT, mBERT and CharacterBERT respectively. For mBERT and CamemBERT, the performance increases when using the last layers, while for CharacterBERT, there is no big difference between the different layers combinations. For example for the setup MODEL+TASK, the accuracy for POS tagging (UPOS) goes from around 70 with the first layer using CamemBERT and mBERT, and reaches 80 using CharacterBERT_arabizi. For CharacterBERT, these scores stay around 80 even when using farther layers (for layers 4 to 7 for instance, the UPOS score is 80.53 and for layers 6-11 it is around 80.34), and while using the last layer gives the best score of 81.19, the latter is still considered around 80 and the stagnation in the scores is hence visible. Contrarily, for CamemBERT and mBERT, the UPOS scores for the MODEL+TASK setup increase from around 70 using only the first layer to above 80 when using layers 6 through 11 (80.39 for CamemBERT and 80.72 for mBERT). The best UPOS score for CamemBERT appears when using the last layer alone (81.14) while for mBERT it is when using layers 6 through 11 (80.72). This clearly illustrates the increase of performance when using higher layers for CamemBERT and mBERT. The same observation can be made for UPOS scores in the MODEL+MLM+TASK setup for mBERT and CamemBERT, and for the other Unlabeled Attachment Score (UAS) and Labeled Attachment Score (LAS) scores as well.

One possible explanation is that the information captured by CharacterBERT layers does not evolve

along the model's layers. The model produces a single embedding for any input token based on an aggregation of the characters embeddings while for BERT-like-models, each sub-word unit in a word is embedded using a WordPiece embedding matrix. Therefore, a possible interpretation is that CharacterBERT learns all the information at the earliest layers as we feed it the whole word directly and not an inconstant count of sub-words when WordPiece vocabulary is in use. Moreover, less than 10% of the 100 000 most frequent sub-words in the *NArabizi* raw data are present in the mBERT vocab due to the high variability of NArabizi.

Layer Config	MODEL+TASK			MODEL+MLM+TASK			Layer Config	MODEL+TASK			MODEL+MLM+TASK		
	UPOS	UAS	LAS	UPOS	UAS	LAS		UPOS	UAS	LAS	UPOS	UAS	LAS
0	68.38	66.21	50.61	70.89	64.74	49.81	0	73.63	65.36	52.08	75.43	66.21	52.98
0-5	77.93	70.51	57.04	80.95	72.21	59.45	0-5	79.40	69.19	57.23	82.37	69.90	58.08
4-7	79.63	70.18	57.56	81.57	72.54	60.92	4-7	80.62	69.61	57.33	83.74	72.45	62.05
6-11	80.39	71.88	58.41	83.65	74.15	62.15	6-11	80.72	68.53	56.76	85.02	72.40	61.63
11	81.14	72.59	60.35	83.08	73.77	62.00	11	80.48	69.19	57.89	84.55	73.82	62.67
all	80.15	70.46	58.08	83.65	74.62	62.62	all	80.25	69.19	56.14	84.31	72.78	61.67

Layer Config	(a) CamemBERT						(b) mBERT		
	NArabizi			Sample OSCAR			NArabizi + Oscar		
	99k			99k			66k+33k		
	UPOS	UAS	LAS	UPOS	UAS	LAS	UPOS	UAS	LAS
0	78.92	70.79	57.99	78.26	69.61	56.76	79.25	69.33	57.14
0-5	79.63	70.23	57.99	79.54	70.13	56.76	79.06	68.05	55.25
4-7	80.53	69.28	57.47	79.77	70.75	57.84	80.10	69.90	56.76
6-11	80.34	70.60	57.84	78.26	68.43	55.58	80.91	69.66	57.80
11	81.19	70.56	58.65	78.83	69.52	56.33	80.67	69.90	57.75
all	79.96	69.47	57.84	80.62	68.76	56.76	80.15	69.90	57.56

Layer Config	UPOS	UAS	LAS
0	78.92	70.79	57.99
0-5	79.63	70.23	57.99
4-7	80.53	69.28	57.47
6-11	80.34	70.60	57.84
11	81.19	70.56	58.65
all	79.96	69.47	57.84

(c) CharacterBERT

Table 9: Performances of the models on the *NArabizi* treebank using **different combinations of the model layers** for embeddings.