

OpenLID-v3: Improving the Precision of Closely Related Language Identification – An Experience Report

Mariia Fedorova Nikolay Arefyev Maja Buljan Jindřich Helcl
Stephan Oepen Egil Rønningstad Yves Scherrer

Language Technology Group
Department of Informatics
University of Oslo
mariiaf@ifi.uio.no

Abstract

Language identification (LID) is an essential step in building high-quality multilingual datasets from web data. Existing LID tools (such as OpenLID or GlotLID) often struggle to identify closely related languages and to distinguish valid natural language from noise, which contaminates language-specific subsets, especially for low-resource languages. In this work we extend the OpenLID classifier by adding more training data, merging problematic language variant clusters, and introducing a special label for marking noise. We call this extended system OpenLID-v3 and evaluate it against GlotLID on multiple benchmarks. During development, we focus on three groups of closely related languages (Bosnian, Croatian, and Serbian; Romance varieties of Northern Italy and Southern France; and Scandinavian languages) and contribute new evaluation datasets where existing ones are inadequate. We find that ensemble approaches improve precision but also substantially reduce coverage for low-resource languages.

1 Introduction

Growing interest in large-scale LLM pre-training data for languages other than English puts a spotlight on robust and broad-coverage language identification (LID). Common pre-training datasets are typically distilled from massive collections of web documents, which are characterized by immense diversity in, for example, genres and domains, degrees of (in)formality, juxtaposition of language and non-language content, presence of machine-generated content, code switching, and other sources of variation. For instance, the two largest and linguistically broadest pre-training datasets, FineWeb 2 (Penedo et al., 2025) and HPLT 3.0 (Oepen et al., 2025), apply LID as a document-level classification task. Specifically, FineWeb builds on the third-party GlotLID classifier (Kargaran et al., 2023), which supports

some 2,000 distinct languages, whereas HPLT has developed a custom classifier dubbed OpenLID (Burchell et al., 2023) covering around 200 languages; see Section 2 for further background. In this work, we seek to shed more light on LID performance and challenges in the realm of noisy web documents, with a particular emphasis on selected groups of closely related languages.

The contributions of this paper are as follows:

- We train a new version of the fully open-source OpenLID system for our experiments and publicly release it as OpenLID-v3.¹ This version covers 194 languages plus ‘not-a-language’ class.
- We evaluate OpenLID-v3 on mainstream LID benchmarks such as FLORES+ and UDHR and show that they are not sufficient for evaluating the quality of similar LID.² In addition, we employ several existing benchmarks for similar languages and create new ones for the BCMS macrolanguage, and for Norwegian Bokmål and Nynorsk.
- We also report negative results on our efforts with two-step coarse-to-fine classification approach, to be found in Appendix F.

2 Related Work

Accurate language identification (LID) is essential for building high-quality multilingual datasets, since documents or segments assigned to an incorrect language can severely contaminate the language-specific subsets, especially for low-resource languages.

¹<https://github.com/hplt-project/openlid>

²<https://github.com/hplt-project/openlid-v3-evaluation>

2.1 LID Tools and Methods

Numerous LID systems have been developed over the years, including `langid.py` (Lui and Baldwin, 2012), Google’s Compact Language Detector 2 (CLD2) and its neural network successor CLD3, the HeLI method (Jauhiainen et al., 2016), LanideNN (Kocmi and Bojar, 2017), or AfroLID (Adebara et al., 2022).

In recent years, classifiers based on the `fastText` model (Joulin et al., 2017) have become the de facto standard in large-scale corpus processing due to their efficiency and strong performance across many languages. `FastText` calculates a hidden representation of a text by summing the embeddings of words and character n-grams from this text. This representation is passed through a linear layer for classification. This approach takes both word- and character-level features into account, making it especially efficient for languages with rich morphology.

In this work, we focus on two `fastText`-based systems: **GlottLID** (Kargaran et al., 2023) and **OpenLID** (Burchell et al., 2023), both designed for multilingual scenarios and widely adopted in massive-scale data curation pipelines such as `FineWeb` (Penedo et al., 2024, 2025) or `HPLT` (de Gibert et al., 2024; Burchell et al., 2025; Oepen et al., 2025). `OpenLID` and `GlottLID` differ primarily in their language coverage and training data selection. `GlottLID` focuses on maximizing the language coverage and supports over 2,000 languages, though it incorporates training data from sources with more restrictive licenses. This extensive coverage results in a larger model size, but makes it suitable for use cases involving low-resource languages. In contrast, `OpenLID` prioritizes fully open-source training data with permissive licenses. The original version (v1) covering 201 languages was later updated (v2) to a reduced coverage of 189 languages. The reduction excluded three problematic languages and consolidated certain language varieties under their macrolanguage labels for compatibility with the `FLORES+` benchmark.

2.2 Broad-Coverage Evaluation Benchmarks

The most widely adopted multilingual language identification benchmarks are `FLORES+` and `UDHR`. **FLORES+**³ is a parallel corpus based on `FLORES-200` (NLLB Team et al., 2022), originally

³https://huggingface.co/datasets/openlanguageata/flores_plus

developed for machine translation evaluation. It features two disjoint publicly available data splits, ‘dev’ and ‘devtest’, and currently covers 225 language varieties. In our experiments, we use the former for validation and the latter for evaluation. The **UDHR** dataset consists of translations of the Universal Declaration of Human Rights across 418 languages,⁴ which provides formal, declarative text in a single domain. The `UDHR` dataset features a test split only.

Additionally, **FastSpell** (Bañón et al., 2024) provides a benchmark specifically designed to discriminate between closely related languages in web documents, which addresses some of the limitations that stem from using relatively clean parallel corpora.

Finally, Oepen et al. (2025) present a manual inspection effort conducted in the context of `HPLT` 3.0 data collection.⁵ In that work, human annotators checked the prediction of `OpenLID-v2` for web documents. They could label documents as containing unnatural language, porn, web artifacts, and incorrect LID. While the data hasn’t been annotated for the correct language, it is possible to construct a LID evaluation dataset from the subset of correctly identified samples. Statistics of this resulting dataset, dubbed **HPLT-LID**, and evaluation results on it are to be found in the Appendix A. We also performed re-annotation of incorrectly classified examples for some languages.

2.3 Discriminating Between Similar Languages

The `VarDial` workshop has a long tradition of shared tasks that focus on language identification in challenging settings, such as the discrimination of closely related languages (e.g. Zampieri et al., 2014; Gaman et al., 2020; Aepli et al., 2023; Chifu et al., 2024). For example, shared tasks have focused on the Bosnian-Croatian-Montenegrin-Serbian macrolanguage or on regional languages of Italy. We leverage datasets made available in this context for a more fine-grained analysis of `OpenLID-v3` (see Section 4).

2.4 Evaluation Metrics

Caswell et al. (2020) argue that the commonly used metric precision (and its derivative, F1-score) are

⁴We use the `udhr-lid` version of the dataset available at <https://huggingface.co/datasets/cis-lmu/udhr-lid>

⁵https://github.com/hplt-project/release3_inspection/tree/main/annot_round1

misleading when evaluating models on the standard LID datasets, as the models face incomparably larger language imbalance in the real web crawls compared to any reasonable labeled dataset. For example, their baseline model achieving a median F1-score of 98% on benchmark data produced a set of monolingual corpora with a median precision of 5% only when evaluated manually on real data. On the other hand, recall and false positive rate (FPR) are not susceptible to class imbalance, and thus are recommended for model comparison.

Since most of our benchmarks are single-label, we report FPR, precision, recall and F1-score for them. For multilabel benchmarks, which are available for some related languages, where short samples may be valid in more than one language, we also report loose (is subset) and exact match metrics (cf. [Fedorova et al., 2025](#)).⁶ These metrics (we refer to them as multilabel classification metrics throughout the paper) are still less strict than our default ones, because they do not depend on the exact number of false positives or negatives in a set of predicted labels.

3 OpenLID-v3

This work is motivated by our experience with OpenLID-v2 in the context of its application to a large-scale web dataset HPLT 3.0, and the results of its manual inspection ([Oepen et al., 2025](#)). We identified a number of issues with OpenLID-v2, which led us to provide an updated version, OpenLID-v3, with a slightly different language inventory. We focus on improving OpenLID-v2 (rather than GlotLID, for example) because of the permissive license of all its training data. The identified issues are as follows:

- Support for the Latin language, which was present in OpenLID-v1, but removed in v2.
- OpenLID-v2 contains only Serbian in Cyrillic script, while Latin script is also widely used in Serbian non-governmental public and private communication. As a result, OpenLID-v2 erroneously classifies Serbian in Latin script as Bosnian or Croatian; according to the manual inspection, half of HPLT 3.0 Bosnian turned out to be Serbian written in Latin script.
- Since OpenLID-v2’s class inventory is limited to 200 languages, any text that is not natural

language (e.g., code, broken encoding – we label these instances as not-a-language, using `zxx_Zxxx` throughout the paper) or a natural language outside of the 200 the model was trained for (we refer to such instances as other), are still predicted to belong to one of the existing classes. In HPLT 3.0, the problem of not-a-language classes was overcome by rule-based filtering after LID, and the problem of the other class was solved by thresholding the predictions by 0.5 softmax scores. However, some classes can still accumulate a large number of not-a-language and other documents. We call this the *trash bin phenomenon*; in HPLT 3.0, Ligurian was found to be such a ‘trash bin’ language ([Oepen et al., 2025](#)).

- Some highly similar languages showed high confusion when tested on the FLORES+ development split and UDHR (Arabic dialects, Persian languages, Bambara and Dyula, Dzongkha and Tibetan, and other languages listed in Table 10).

Finally, the amount of identified data for certain languages in the HPLT dataset is very small compared to its number of speakers, e.g. Bengali and Tamil. See Table 10 for the full list of such languages.

When developing OpenLID-v3, we approached these issues in the following ways:

- We merged 8 Arabic dialects into one Arabic macrolanguage: `ara_Arab`; and two Persian varieties: `pes_Arab` and `prs_Arab` into the Farsi macrolanguage `fas_Arab`. This follows the OpenLID-v1 approach. Likewise, we merged Bambara and Dyula, which are mutually intelligible ([Amuzu and Singler, 2014](#)).
- We introduced the not-a-language class `zxx_Zxxx`, using the GlotLID training data labeled as `und_*` and `zxx_*`.
- We extended the training data for several languages, including Latin, and Serbian written in Latin script. We relied on the subsets of the GlotLID training data which were reported by its authors to not be noisy. We also removed one possibly noisy OpenLID-v2 training subset and added the most recent Wikipedia dumps, where adding GlotLID training data showed insufficient improvements during validation. The affected lan-

⁶<https://github.com/ltgoslo/slide>

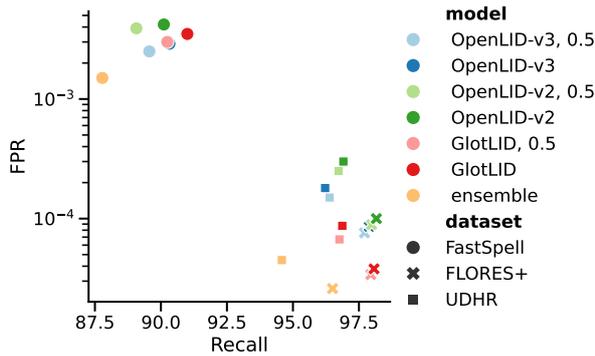


Figure 1: Model comparison on three LID benchmarks, FPR (lower is better) vs. recall (higher is better). For the full set of metrics see Table 9. FLORES+ refers to the devtest split, and FastSpell refers to our version, which excludes Nynorsk.

guages and data sources used are summarized in Table 10 in the Appendix.

- We also experimented with ensembling OpenLID-v3 and GlotLID by top-1 and top-3 agreement. Top-3 agreement worsened the results, which is expected for a single-label model, where only one class gets a relatively high softmax score, and all other classes are a long tail of small numbers.

The list of language labels supported by OpenLID-v3 is given in Table 15 in Appendix G.

3.1 Results

Figure 1 compares the results of our OpenLID-v3 model with the previous version (OpenLID-v2) as well as with GlotLID, focusing on the popular LID benchmarks: FLORES+, UDHR, and FastSpell. Following HPLT 3.0, we experimented with applying the same softmax thresholding of 0.5 to all three models.

OpenLID-v3 is on par with both OpenLID-v2 and GlotLID; the difference with and without thresholding softmax scores is not large among models, but the results are slightly better for OpenLID-v3 with thresholding, and for GlotLID without. This different behavior may be explained by the fact that GlotLID’s ‘trash bins’ fall outside of the classes we evaluate on. In our further experiments, we use OpenLID-v3 with thresholding, and GlotLID without, unless otherwise specified, including the ensembling approach.

Top-1 ensembling resulted in the lowest FPR across all datasets. Based on our experiments, this

approach was used to produce HPLT 4.0 datasets.⁷

4 Case Studies on Related languages

Results of both OpenLID and GlotLID are high on multilingual benchmarks. However, lower precision for similar languages may be hidden when averaging their scores with those of languages that are easy to distinguish. For this reason, we perform additional evaluation for groups of closely related languages. We choose three groups of such languages, within those with updated training data, based on benchmark availability and our language expertise.

4.1 Bosnian, Croatian, Montenegrin and Serbian

Bosnian, Croatian and Serbian are, along with Montenegrin, South Slavic languages spoken in the West Balkans; part of a group of historically and geographically close languages that are, for the most part, mutually intelligible to speakers of any one of them. They differ by variations in vocabulary and grammatical features – most notably, the orthography of the Slavic *jat* (ě) (Mešanović-Meša, 2011; Karavdić, 2017). Montenegrin is only present in GlotLID class inventory; for this reason, we do not focus on it. The analysis of its GlotLID predictions is to be found in Appendix C.

Test data. We use three datasets for the BCMS language group.⁸ First, BCMS **Twitter user** dataset (Ljubešić and Rupnik, 2022b), also featured in the BENCHiC BCMS benchmark (Rupnik et al., 2023). We use the version with multi-label annotations (Miletić and Miletić, 2024), comprising roughly 123 instances (in the test set), where each instance corresponds to one user and contains many texts written by that user. Second, we use the BCS portion of **ParlaSent** (Mochtak et al., 2023), comprising roughly 18,000 sentences from transcriptions of parliamentary debates, annotated with the speaker’s country of origin. Third, we make use of the **HPLT-LID** data. Since Serbian in Latin script was absent in OpenLID-v2, pre-annotation was only performed for Bosnian and Croatian. In total, 804 sentences were pre-annotated (402 predicted to be Bosnian and 402 predicted to be Croatian). Out of them, 114 Bosnian samples and 13 Croatian were pre-annotated to be false positives.

⁷<https://hplt-project.org/datasets/v4.0>

⁸Exact scores on these benchmarks are to be found in the Appendix C.

Model	Bosnian			Croatian			Serbian Latin		
	FPR	Prec.	Rec.	FPR	Prec.	Rec.	FPR	Prec.	Rec.
<i>Twitter users</i>	n=21			n=22			n=94		
OpenLID-v3	0.23	37.21	76.19	0.026	80.00	54.54	0	100.00	30.85
OpenLID-v2	0.8	16.96	90.48	0.017	84.62	50.00	0	0	0
GlottLID	0.42	16.95	47.62	0.26	36.17	77.27	0.12	82.14	24.49
Ensemble	0	0	0	0	0	0	0	0	0
<i>ParlaSent</i>	n=153			n=1,387			n=1,060		
OpenLID-v3	0.22	11.59	47.71	0.13	87.49	79.16	0.019	95.00	53.77
OpenLID-v2	0.44	0.08	57.51	0.26	78.05	80.25	0	0	0
GlottLID	0.26	9.59	43.79	0.242	81.02	90.48	0.006	97.10	31.60
Ensemble	0.14	13.15	34.64	0.11	89.25	77.21	0.003	98.38	28.77
<i>HPLT-LID (reannotated)</i>	n=5			n=7			n=110		
OpenLID-v3	0.15	5.26	40.00	0.03	42.86	85.71	0	100.00	50.00
OpenLID-v2	0.9	1.85	40.00	0.09	21.43	42.86	0	0	0
GlottLID	0.4	0.01	0.2	0.05	31.58	85.71	0	100.00	3.64
Ensemble	0.146	2.70	20.00	0.03	46.15	85.71	0	100.00	3.64

Table 1: Performance comparison across three benchmark datasets showing false positive rate (FPR), precision (Prec.), and recall (Rec.) for Bosnian, Croatian, and Serbian in Latin script. OpenLID-v3 uses softmax threshold 0.5. Ensemble represents top-1 agreement between OpenLID-v3 and GlotLID. Best values per metric and language are in **bold**.

We performed re-annotation of these false positive samples to their correct labels; the majority of them was found to be Serbian. The annotation was done by a native speaker of Croatian with linguistic background. One sample was detected to be valid both in Croatian and Serbian and annotated as multilabel; we excluded it from metric calculation.

4.1.1 Quantitative Evaluation

Table 1 presents FPR, precision and recall per language on Twitter user, ParlaSent and reannotated BCS part of HPLT-LID.

None of the models scored high on all three datasets. As HPLT 4.0 developers, we were mostly interested in results on Twitter, since its texts were the closest to the noisy web data, containing hyperlinks, emodjis etc. OpenLID-v3 turned out to be the best for Bosnian and Serbian on it; slightly higher precision of Croatian obtained from OpenLID-v2 comes at cost of lack of Serbian class label, which is not what we aim at. Importantly, GlotLID and OpenLID-v3 always disagree on Twitter data, which emphasizes that ensembling should be done cautiously for particular language groups. For the future work on improving predictions for the Twitter dataset, there is strong need for multilabel training data, silver approach from (Fedorova et al., 2025) might be helpful for future work.

Scores on ParlaSent were the highest, especially

Model	Loose Acc.	Exact Acc.	bos F1	hrv F1	srp F1
OpenLID-v3	46.34	40.65	57.14	72.73	49.15
GlottLID	40.65	33.33	28.17	55.74	40.00

Table 2: Loose accuracy, exact match accuracy and loose F1s per languages on multilabel BCS test data (Twitter). Ensemble@1 never agrees (always other, all metrics zero).

in precision of Serbian, which proves it to be the ‘easiest’ dataset.

Results of evaluation on reannotated HPLT-LID were similar to those on Twitter: OpenLID-v3 performed the best, and there were a high disagreement between it and GlotLID. However, the confusion of Bosnian and Serbian was still high.

Table 2 shows multilabel classification metrics on Twitter data. By these metrics, OpenLID-v3 is a clear winner over GlotLID.

4.1.2 Common Errors

We performed manual analysis of BCS predictions from all models and observed the following seven most frequent error patterns:

NE confusion Named entities are a common source of confusion. The presence of a different country name (e.g. Serbian news article about Croa-

(gold/predicted label)	b/h	b/s	b/x	h/b	h/s	h/x	s/b	s/h	s/x	% of total
NE confusion	12.0	-	2.0	4.0	-	2.0	10.0	6.0	8.0	6.1
lexical overlap	34.0	-	20.0	26.0	-	2.0	44.0	12.0	6.0	20.5
historic forms	-	-	-	36.0	-	2.0	-	-	-	5.8
<i>da</i> confusion	-	66.6	4.0	-	100.0	-	-	-	-	1.7
ungrammatical syntax	10.0	-	32.0	10.0	-	64.0	6.0	14.0	28.0	22.5
total ambiguity	26.0	-	12.0	6.0	-	20.0	-	34.0	26.0	19.9
misabeled minority rep	-	33.3	4.0	2.0	-	-	-	-	-	1.2
(other/unknown)	18.0	-	4.0	16.0	-	10.0	40.0	34.0	32.0	22.2

Table 3: V3-ensemble evaluation on ParlaSent; percentage of error occurrences, per sample of 50 mislabeled documents for each mismatched pair of labels. Exceptions are b/s, b/x, and h/s, which have a total of 2, 40, and 3 datapoints each, respectively. The bottom row, *other*, is a catch-all for unambiguous documents with clear language markers, which were nevertheless mislabeled.

tia), an individual or institution that adheres to the other language’s naming convention, or even word type confusion (“*Tome se niko ne raduje*” (“Nobody’s looking forward to that”, compared with the common Slovenian name *Tome*) identified as Slovenian), leads to mislabelling, due to lack of other language markers.

Lexical overlap For a human reader, the current orthography of *jat* is a strong discriminator between Serbian and Bosnian or Croatian, the former preferring the *e* form in the standard language (e.g. *videti, mleko*), and the other two using *je/ije* (*vidjeti, mlijeko*). However, documents are frequently mislabeled between Bosnian and Serbian, despite the presence of clear *e/ije* indicators, if the document also contains roots and noun/adjective inflections that are common to both Bosnian and Serbian, whereas Croatian uses a different surface form of the word (e.g. *obavezno* (Bosnian, Serbian) vs. *obvezatno* (Croatian)). The frequency of these error occurrences (*lexical overlap* in 3) seems to indicate that common roots and shared lexemes are a stronger signal than *jat* orthography.

Historic forms Similarly, all models frequently mislabel ParlaSent documents in Croatian as either Bosnian or Serbian, when the speakers slip into historic lexeme forms that were part of the shared Serbo-Croatian language, but are not present in current standard Croatian. Older speakers and colloquial language still frequently use *historic forms*, which again leads to mislabeling despite the unambiguous grammatical markers.

***da* confusion** Another theoretically strong indicator is the difference between the future tense in Croatian (modal verb + infinitive, e.g. *(ho)ću glasati, glasat ću*) and Serbian (modal verb + *da* + present simple, e.g. *(ho)ću da glasam*); both forms

are accepted in Bosnian. Frequently, Croatian documents are mislabelled as Serbian when there are multiple occurrences of the conjunction *da* (meaning *that*; not to be confused with the particle *da* (yes)). E.g., “*Ne sumnjam da je ovaj zakon nekima donio dobro i ne sumnjam da oni hvale ovaj zakon.*” (“I do not doubt **that** this legislation brought good to some, and I do not doubt **that** they praise it.”) – Croatian mislabelled as Serbian – vs. “*Neću da glasam za taj zakon.*” (“**I will not / don’t want to vote** for this legislation.”) – true Serbian.

Ungrammatical syntax Ungrammatical syntax comprises documents in non-standard language; mainly long run-on sentences, with a lack of punctuation, incorrect syntax, frequent interjections, and generally language that is highly colloquial in structure and word choice. While this mislabeling occurs across most pairings, these documents are most frequently labeled as *not-a-language* (zxx_Zxxx), particularly when there is a lack of any distinguishing grammatical or lexical language markers.

Finally, some datapoints are mislabeled due to extralinguistic factors.

Total ambiguity Some documents have no clear language markers between BCS. Most of these documents are mislabeled across models, and would arguably be hard even for a human annotator and native speaker to tease apart, due to a lack of unambiguous linguistic features.

Mislabeled minority representative Mislabeled minority representative is specific to the ParlaSent dataset, where statements are labeled according to the national parliament in which the discussion occurred. In a handful of cases, it is clear that the speaker is speaking a minority language, and the

model has identified it correctly, though it doesn't match the parliament country of origin.

We performed a quantitative study of these error types on the ParlaSent predictions obtained from the best model (OpenLID-v3 top-1 ensemble with GlotLID). The results are presented in the Table 3. The last row (*other/unknown*) quantifies unambiguous documents with clear language markers, which were nevertheless mislabeled by the model. This is a particularly frequent occurrence for Latin-script Serbian.

4.2 Romance Languages of Italy and France

The HPLT 3.0 compilation efforts reported language identification issues connected to Ligurian. 31% of the samples annotated as Occitan in the UDHR dataset were predicted as Ligurian by OpenLID-v2. Upon closer inspection, it turned out that the Occitan part of UDHR contained seven translations – three truly Occitan ones and four Francoprovençal ones (see Table 12 in the Appendix). Although Ligurian, Occitan and Francoprovençal are all Romance languages and are spoken in adjacent areas (Southern France, Northwestern Italy and Western Switzerland), they are generally considered to belong to different genealogical sub-groups (Ledgeway and Maiden, 2016; Ramponi, 2024). Since Francoprovençal was missing from OpenLID-v2's class inventory, its instances were erroneously predicted as Ligurian.

This suggests that (a) multilingual benchmarks cannot be fully relied on when it comes to closely related, unstandardized languages, and (b) LID tools struggle to identify low-resource languages especially if higher-resourced closely related varieties are present in their training data.

We further investigate the performance of OpenLID-v3 on these Romance languages. As we were unable to find a test set that covered Francoprovençal, Ligurian and Occitan,⁹ we use the ITDI dataset instead, which covers several languages and dialects of Italy (Aepli et al., 2022). The metrics were only calculated for Ligurian, Venetian and Friulian, as only these three languages are predicted by both OpenLID-v3 and GlotLID.

The results are presented in Table 4. While GlotLID is better than OpenLID-v3 both in precision and recall, ensembling them achieves the best precision and the lowest FPR. For both mod-

⁹While there exists an effort on Occitan LID (Miletic and Scherrer, 2022), we could not use it in our experiments, as it contains silver labels only.

Model	FPR	F1	Precision	Recall
OpenLID-v3	0.01	80.88	96.51	74.22
GlotLID	0.01	82.37	97.60	75.15
ensemble	0.004	76.68	98.47	68.95

Table 4: Comparison of OpenLID-v3 and GlotLID on the ITDI test set restricted to Ligurian, Venetian and Friulian (4,744 samples).

els, Venetian confusion is the highest with Italian, but OpenLID-v3 predicts fewer Italian false positives and more not-a-language ones (which is never the case for GlotLID). Friulian is mostly predicted correctly, more often with OpenLID-v3 (1,284 true positives) than with GlotLID (1,275). Ligurian is also more often predicted correctly by both models, with Sicilian and Friulian as the top confusion instances. For OpenLID-v3, the third most frequently predicted class is not-a-language. Ensembling the two models removes 45% of Venetian samples, while the number of Ligurian and Friulian ones remains roughly the same.

In sum, this case study emphasizes the importance of fine-grained benchmarks for related Romance languages. GlotLID might be a good choice as a standalone model; however, the most precise results are obtained when it is ensembled with OpenLID-v3, although at the price of removing samples from other easily confused languages such as Venetian.

4.3 Scandinavian Languages

Lastly, we focus on the North Germanic languages of Mainland Scandinavia, i.e. Norwegian Bokmål, Norwegian Nynorsk, Danish and Swedish. In addition to FLORES+ and UDHR, which cover all four languages, we experiment with two language-group-specific datasets: First, we use **SLIDE** (Fedorova et al., 2025), a multilabel dataset based on clean instances from Universal Dependencies (Nivre et al., 2020) (6,950 samples in total), containing many other languages. Second, we use **Nordic DSL** (Haas and Derczynski, 2021), a single-label dataset containing noisy sentences sourced from Wikipedia and featuring Faroese and Icelandic, in addition to the four aforementioned languages.

We did not use the **FastSpell** dataset for the calculation of metrics, as upon manual examination, we noticed a large share of its Nynorsk samples to be valid in both Bokmål and Nynorsk. We re-

Model	Norwegian Bokmål			Danish			Norwegian Nynorsk			Swedish			Other		
	FPR	Prec.	Rec.	FPR	Prec.	Rec.	FPR	Prec.	Rec.	FPR	Prec.	Rec.	FPR	Prec.	Rec.
<i>SLIDE</i>	n=2,098			n=677			n=1,628			n=1,250			n=1,745		
OpenLID-v3	0.04	88.73	87.04	0.02	81.81	83.75	0.03	89.66	84.64	0.008	96.11	96.80	0.03	89.68	95.13
OpenLID-v2	0.05	87.33	88.04	0.01	85.27	80.35	0.03	90.47	84.58	0.008	95.88	96.8	0.03	89.63	95.58
GlottLID	0.047	88.32	90.47	0.02	80.33	86.26	0.008	96.57	81.33	0.009	95.83	97.44	0.03	92.07	99.20
Ensemble	0.027	92.56	84.84	0.01	88.85	81.24	0.006	97.37	79.55	0.003	98.29	96.8	0.1	75.86	99.77
<i>Nordic DSL</i>	n=14,960			n=14,960			n=14,960			n=14,960			n=14,960		
OpenLID-v3	0.007	96.04	83.22	0.012	94.10	93.56	0.017	91.98	96.87	0.002	99.07	93.33	0.039	92.72	99.44
OpenLID-v2	0.02	91.85	85.37	0.007	96.12	89.78	0.02	91.97	96.75	0.002	98.77	92.73	0.04	92.82	99.56
GlottLID	0.01	94.58	94.50	0.006	97.09	93.77	0.006	97.20	96.47	0.004	98.16	93.69	0.025	95.16	99.36
Ensemble	0.005	97.19	82.78	0.003	98.17	91.82	0.004	98.13	95.92	0.001	99.35	92.01	0.079	86.32	99.74
<i>FLORES+</i>	n=1,012			n=1,012			n=1,012			n=1,012			n=212,520		
OpenLID-v3	1e-4	97.61	96.84	3e-5	99.31	99.21	7e-5	98.42	98.22	0	1.0	99.90	3e-3	99.99	1.0
OpenLID-v2	1e-4	96.84	96.84	2e-5	99.60	98.12	9e-5	98.12	98.22	0	1.0	99.90	0.004	99.99	99.99
GlottLID	1e-4	97.10	99.41	1e-5	99.70	99.51	1e-5	99.67	97.53	0	1.0	1.0	0	1.0	1.0
Ensemble	7e-5	98.39	96.74	6e-6	99.80	98.91	1e-5	99.70	97.23	0	1.0	99.90	1e-2	99.98	1.0
<i>UDHR</i>	n=62			n=61			n=58			n=61			n=27,515		
OpenLID-v3	4e-5	98.39	98.39	4e-5	98.36	98.36	8e-5	96.61	98.28	7e-5	96.83	1.0	8e-3	99.99	99.98
OpenLID-v2	7e-5	96.88	1.0	7e-5	96.77	98.36	4e-5	98.28	98.28	1e-4	95.31	1.0	0.004	99.99	99.97
GlottLID	4e-5	98.41	1.0	1e-4	93.75	98.36	4e-5	98.31	1.0	4e-5	98.39	1.0	0	1.0	99.98
Ensemble	4e-5	98.39	98.39	0	1.0	98.36	0	1.0	98.28	0	1.0	1.0	8e-3	99.99	1.0

Table 5: Performance comparison across four benchmark datasets showing the false positive rate (FPR), precision (Prec.), and recall (Rec.) for Scandinavian and other languages. *OpenLID-v3* uses softmax threshold 0.5. *Ensemble* represents top-1 agreement between *OpenLID-v3* and *GlottLID*. Best values per metric and language are in **bold**.

annotated this subset and evaluated it separately.

Table 5 summarizes the results. We find that GlotLID obtains the best recall, while OpenLID shows better false positive rate and precision. Precision and FPR are further improved by ensembling GlotLID and OpenLID-v3. OpenLID-v3 performs on par or slightly better than OpenLID-v2 on all benchmarks. Furthermore, the confusion matrices (available in the repository¹⁰) show that the two Norwegian varieties, Bokmål and Nynorsk, are most easily confused. Appendix E provides additional results: Table 14 shows the effect of applying softmax threshold across all models and datasets, and Table 13 shows the effect of ensembling across datasets. Both tables report multilabel classification metrics.

HPLT-LID The HPLT 3.0 manual inspection data (Oepen et al., 2025) only covers one Scandinavian language, namely Norwegian Bokmål. 402 documents were annotated as being correctly predicted Bokmål or not. We process this annotation in the same way as we did for BCMS annotation, excluding porn, unnatural language and artifacts. Documents with correctly identified language formed the dataset of documents guaranteed

¹⁰https://github.com/ltgoslo/slide/tree/main/src/eval_logs/wardial_2026

to be Bokmål, with 304 samples in total. 5 documents were annotated to be in a language other than Bokmål. We manually reviewed predictions of these 5 false positive Bokmål documents. Both OpenLID-v3 and GlotLID still predict them to be Bokmål. 3 of them are non-fluent Bokmål documents with many grammatical mistakes (articles of the wrong gender, articles and verbs omitted), as if written by foreign speakers, and 2 include code-switching, with Bokmål still being the predominant language.

FastSpell. A native speaker of Norwegian relabeled the FastSpell Nynorsk subset in a multilabel way, finding 40% of samples to be (also) valid Bokmål. The annotator also reported that many word sequences were not full sentences and some texts were just menu choices from a web page. The evaluation of OpenLID-v3 and GlotLID on the resulting dataset is presented in Table 6. While both models performed poorly in Bokmål detection, OpenLID-v3 performed slightly better for Nynorsk and produced fewer Danish and Swedish false positives, but more other false positives than GlotLID.

The manual analysis of the predictions showed that errors are often caused by named entities valid in any language using Latin script (“Georg Johannes Toft”), dates (“Levering 15. juli 2018”)

Model	NB F1	NN F1	DA FP	SV FP	Other FP
OpenLID-v3	71.90	92.19	16	8	52
GlottLID	79.07	91.14	27	19	16

Table 6: GlotLID (no threshold) and OpenLID-v3 (softmax threshold 0.5) on the re-annotated FastSpell Nynorsk data. F1 is for loose F1, FP is for false positives. 406 Nynorsk samples, 163 Bokmål samples.

and foreign words (“tomter i Parque natural cabo de gata”). These findings correspond to the error analysis from Fedorova et al. (2025) and further prove that shorter sequences, especially from the web, are the hardest for LID models.

In conclusion, we recommend OpenLID-v3 for general Scandinavian language identification if precision is prioritized, and GlotLID if recall is prioritized. For discriminating between Bokmål and Nynorsk, it is recommended to use a LID model trained specifically to separate these languages.

5 Conclusion

We have evaluated the performance of three LID models (two OpenLID versions and GlotLID) on four multilingual benchmarks (two based on ‘clean’ text and two based on web documents). Our new model, OpenLID-v3, performs on par or better in precision compared to its predecessors OpenLID-v2 and GlotLID, while the best precision is achieved by top-1 ensembling with GlotLID. We have also shown that reliable evaluation of LID models including similar languages should be done on benchmarks specific to these languages, even if overall metrics on large multilingual benchmarks are high. Since web data may contain short texts valid in more than one language, there is a need for more multilabel training data and benchmarks.

Limitations

The first limitation that we acknowledge is data mismatch between the evaluation and the intended use – we would ideally evaluate on web text. However, large-scale web LID data of sufficient quality were not available at the time of conducting our experiment. We are aware of the initiative by Suarez et al. (2026), but leave evaluation on their dataset for future work.

We have done our best to control for data contamination, where there was a known intersection of training and test data. However, this was not possible for tests on Nordic DSL, as this dataset,

while consisting of Wikipedia data, was heavily preprocessed by its authors.

A common limitation of all models under evaluation is the unavailability of fully parallel datasets for all languages of interest, which might cause models to overfit to some concepts and named entities that are more common in certain languages.

Ethical Considerations

All the new data annotations were done by the authors voluntarily and without any monetary compensation.

We did not check any of OpenLID-v3’s training data for inappropriate or biased content. We believe it does not do much harm, since the model is not generative. However, we can imagine some model predictions to be biased, particularly if used not as a classifier, but as a source of semantic representations.

A large amount of the current research is focused on collecting data for training large instruction-tuned generative models, capable of outputting grammatically and orthographically correct standard language, because of the high demand for such models in the public sector. This may result in loss of data written in low-resource varieties of mainstream languages, along with possibly valuable cultural knowledge.

Acknowledgements

This project was supported by the European Union Horizon Europe project no. 101070350 (HPLT). The computations were performed on resources provided by Sigma2 – the National Infrastructure for High-Performance Computing and Data Storage in Norway.

References

- Ife Adebara, AbdelRahim Elmadany, Muhammad Abdul-Mageed, and Alcides Inciarte. 2022. *AfroLID: A neural language identification tool for African languages*. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1958–1981, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Noëmi Aepli, Antonios Anastasopoulos, Adrian Chifu, William Domingues, Fahim Faisal, Mihaela Găman, Radu Tudor Ionescu, and Yves Scherrer. 2022. Findings of the VarDial evaluation campaign 2022. In *Proceedings of the Ninth Workshop on NLP for Similar Languages, Varieties and Dialects*, Gyeongju,

- Republic of Korea. International Committee on Computational Linguistics (ICCL).
- Noëmi Aepli, Çağrı Çöltekin, Rob Van Der Goot, Tommi Jauhiainen, Mourhaf Kazzaz, Nikola Ljubešić, Kai North, Barbara Plank, Yves Scherrer, and Marcos Zampieri. 2023. [Findings of the VarDial evaluation campaign 2023](#). In *Tenth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial 2023)*, pages 251–261, Dubrovnik, Croatia. Association for Computational Linguistics.
- Milind Agarwal, Md Mahfuz Ibn Alam, and Antonios Anastasopoulos. 2023. [LIMIT: Language identification, misidentification, and translation using hierarchical models in 350+ languages](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14496–14519, Singapore. Association for Computational Linguistics.
- Evershed Amuzu and John Singler. 2014. [Codeswitching in west africa](#). *International Journal of Bilingualism*, 18:329–345.
- Marta Bañón, Gema Ramírez-Sánchez, Jaume Zaragoza-Bernabeu, and Sergio Ortiz Rojas. 2024. [FastSpell: The LangID magic spell](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 7133–7140, Torino, Italia. ELRA and ICCL.
- Laurie Burchell, Alexandra Birch, Nikolay Bogoychev, and Kenneth Heafield. 2023. [An open dataset and model for language identification](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 865–879, Toronto, Canada. Association for Computational Linguistics.
- Laurie Burchell, Ona de Gibert, Nikolay Arefyev, Mikko Aulamo, Marta Bañón, Pinzhen Chen, Mariia Fedorova, Liane Guillou, Barry Haddow, Jan Hajič, Jindřich Helcl, Erik Henriksson, Mateusz Klimaszewski, Ville Komulainen, Andrey Kutuzov, Joonas Kytöniemi, Veronika Laippala, Petter Mæhlum, Bhavitvya Malik, and 16 others. 2025. [An expanded massive multilingual dataset for high-performance language technologies \(HPLT\)](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 17452–17485, Vienna, Austria. Association for Computational Linguistics.
- Isaac Caswell, Theresa Breiner, Daan van Esch, and Ankur Bapna. 2020. [Language ID in the wild: Unexpected challenges on the path to a thousand-language web text corpus](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6588–6608, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Adrian-Gabriel Chifu, Goran Glavaš, Radu Tudor Ionescu, Nikola Ljubešić, Aleksandra Miletić, Filip Miletić, Yves Scherrer, and Ivan Vulić. 2024. [VarDial evaluation campaign 2024: Commonsense reasoning in dialects and multi-label similar language identification](#). In *Proceedings of the Eleventh Workshop on NLP for Similar Languages, Varieties, and Dialects (VarDial 2024)*, pages 1–15, Mexico City, Mexico. Association for Computational Linguistics.
- Ona de Gibert, Graeme Nail, Nikolay Arefyev, Marta Bañón, Jelmer van der Linde, Shaoxiong Ji, Jaume Zaragoza-Bernabeu, Mikko Aulamo, Gema Ramírez-Sánchez, Andrey Kutuzov, Sampo Pyysalo, Stephan Oepen, and Jörg Tiedemann. 2024. [A new massive multilingual dataset for high-performance language technologies](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 1116–1128, Torino, Italia. ELRA and ICCL.
- Mariia Fedorova, Jonas Sebulon Frydenberg, Victoria Handford, Victoria Ovedie Chruickshank Langø, Solveig Helene Willoch, Marthe Løken Midtgaard, Yves Scherrer, Petter Mæhlum, and David Samuel. 2025. [Multi-label Scandinavian language identification \(SLIDE\)](#). In *Proceedings of the Third Workshop on Resources and Representations for Under-Resourced Languages and Domains (RESOURCEFUL-2025)*, pages 179–189, Tallinn, Estonia. University of Tartu Library, Estonia.
- Mihaela Gaman, Dirk Hovy, Radu Tudor Ionescu, Heidi Jauhiainen, Tommi Jauhiainen, Krister Lindén, Nikola Ljubešić, Niko Partanen, Christoph Purschke, Yves Scherrer, and Marcos Zampieri. 2020. [A report on the VarDial evaluation campaign 2020](#). In *Proceedings of the 7th Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 1–14, Barcelona, Spain (Online). International Committee on Computational Linguistics (ICCL).
- Cyril Goutte, Serge Léger, and Marine Carpuat. 2014. [The NRC system for discriminating similar languages](#). In *Proceedings of the First Workshop on Applying NLP Tools to Similar Languages, Varieties and Dialects*, pages 139–145, Dublin, Ireland. Association for Computational Linguistics and Dublin City University.
- René Haas and Leon Derczynski. 2021. [Discriminating between similar Nordic languages](#). In *Proceedings of the Eighth Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 67–75, Kiyv, Ukraine. Association for Computational Linguistics.
- Tommi Jauhiainen, Krister Lindén, and Heidi Jauhiainen. 2016. [HeLI, a word-based backoff method for language identification](#). In *Proceedings of the Third Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial3)*, pages 153–162, Osaka, Japan. The COLING 2016 Organizing Committee.
- Tommi Jauhiainen, Marco Lui, Marcos Zampieri, Timothy Baldwin, and Krister Lindén. 2019. Automatic language identification in texts: A survey. *Journal of Artificial Intelligence Research*, 65:675–782.

- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. [Bag of tricks for efficient text classification](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431, Valencia, Spain. Association for Computational Linguistics.
- Zenaida Karavdić. 2017. Kako prepoznati bosanski, crnogorski, hrvatski i srpski standardni jezik? *BOSNIACA-časopis Nacionalne i univerzitetske biblioteke Bosne i Hercegovine*, (22):34–46.
- Amir Hossein Kargaran, Ayyoob Imani, François Yvon, and Hinrich Schuetze. 2023. [GlotLID: Language identification for low-resource languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 6155–6218, Singapore. Association for Computational Linguistics.
- Tom Kocmi and Ondřej Bojar. 2017. [LanideNN: Multilingual language identification on character window](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 927–936, Valencia, Spain. Association for Computational Linguistics.
- Adam Ledgeway and Martin Maiden. 2016. *The Oxford Guide to the Romance Languages*. Oxford University Press.
- Nikola Ljubešić and Peter Rupnik. 2022a. [The news dataset for discriminating between bosnian, croatian and serbian SETimes.HBS 1.0](#). Slovenian language resource repository CLARIN.SI.
- Nikola Ljubešić and Peter Rupnik. 2022b. [The twitter user dataset for discriminating between bosnian, croatian, montenegrin and serbian twitter-HBS 1.0](#). Slovenian language resource repository CLARIN.SI.
- Marco Lui and Timothy Baldwin. 2012. [langid.py: An off-the-shelf language identification tool](#). In *Proceedings of the ACL 2012 System Demonstrations*, pages 25–30, Jeju Island, Korea. Association for Computational Linguistics.
- Emira Mešanović-Meša. 2011. *Kontrastivna analiza bosanskog, hrvatskog i srpskog jezika u zakonima Federacije Bosne i Hercegovine*. Slavistički komitet BiH.
- Aleksandra Miletić and Filip Miletić. 2024. A gold standard with silver linings: Scaling up annotation for distinguishing bosnian, croatian, montenegrin and serbian. In *Proceedings of the Fourth Workshop on Human Evaluation of NLP Systems (HumEval)@ LREC-COLING 2024*, pages 36–46.
- Aleksandra Miletic and Yves Scherrer. 2022. [OcWikiDisc: a corpus of Wikipedia talk pages in Occitan](#). In *Proceedings of the Ninth Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 70–79, Gyeongju, Republic of Korea. Association for Computational Linguistics.
- Michal Mochtak, Peter Rupnik, and Nikola Ljubešić. 2023. [The parlament multilingual training dataset for sentiment identification in parliamentary proceedings](#). (arXiv:2309.09783). ArXiv:2309.09783 [cs].
- 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 Twelfth Language Resources and Evaluation Conference*, pages 4034–4043, Marseille, France. European Language Resources Association.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, and 20 others. 2022. [No language left behind: Scaling human-centered machine translation](#). *Preprint*, arXiv:2207.04672.
- Stephan Oepen, Nikolay Arefev, Mikko Aulamo, Marta Bañón, Maja Buljan, Laurie Burchell, Lucas Charpentier, Pinzhen Chen, Mariia Fedorova, Ona de Gibert, Barry Haddow, Jan Hajič, Jindřich Helcl, Andrey Kutuzov, Veronika Laippala, Zihao Li, Risto Luukkonen, Bhavitvya Malik, Vladislav Mikhailov, and 13 others. 2025. [HPLT 3.0: Very large-scale multilingual resources for LLM and MT. mono- and bilingual data, multilingual evaluation, and pre-trained models](#). *Preprint*, arXiv:2511.01066.
- Guilherme Penedo, Hynek Kydlíček, Loubna Ben alal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. 2024. [The fineweb datasets: Decanting the web for the finest text data at scale](#). In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Guilherme Penedo, Hynek Kydlíček, Vinko Sabolčec, Bettina Messmer, Negar Foroutan, Amir Hossein Kargaran, Colin Raffel, Martin Jaggi, Leandro Von Werra, and Thomas Wolf. 2025. [Fineweb2: One pipeline to scale them all – adapting pre-training data processing to every language](#). *Preprint*, arXiv:2506.20920.
- Alan Ramponi. 2024. [Language varieties of italy: Technology challenges and opportunities](#). *Transactions of the Association for Computational Linguistics*, 12:19–38.
- Daniel Romić. 2024. [Heritage bosnian, croatian, and serbian spoken by second generation speakers in germany he-bcs-ge](#).
- Peter Rupnik, Taja Kuzman, and Nikola Ljubešić. 2023. [BENCHi-lang: A benchmark for discriminating between Bosnian, Croatian, Montenegrin and Serbian](#). In *Tenth Workshop on NLP for Similar Languages*,

Varieties and Dialects (VarDial 2023), pages 113–120, Dubrovnik, Croatia. Association for Computational Linguistics.

Pedro Ortiz Suarez, Laurie Burchell, Catherine Arnett, Rafael Mosquera-Gómez, Sara Hincapie-Monsalve, Thom Vaughan, Damian Stewart, Malte Ostendorff, Idris Abdulmumin, Vukosi Marivate, Shamsuddeen Hassan Muhammad, Atnafu Lambebo Tonja, Hend Al-Khalifa, Nadia Ghezaiel Hammouda, Verah Otiende, Tack Hwa Wong, Jakhongir Saydaliev, Melika Nobakhtian, Muhammad Ravi Shulthan Habibi, and 78 others. 2026. [Commonlid: Re-evaluating state-of-the-art language identification performance on web data](#). *Preprint*, arXiv:2601.18026.

Marcos Zampieri, Liling Tan, Nikola Ljubešić, and Jörg Tiedemann. 2014. [A report on the DSL shared task 2014](#). In *Proceedings of the First Workshop on Applying NLP Tools to Similar Languages, Varieties and Dialects*, pages 58–67, Dublin, Ireland. Association for Computational Linguistics and Dublin City University.

A Multilingual Benchmarks

In all evaluations on all datasets, metrics were only calculated on subset of languages, present in all models being compared. For GlotLID, we always used its the most recent version at the moment of writing, GlotLID-v3.

Table 7 shows statistics of the HPLT-LID dataset. Only those samples were accepted, which were not unnatural language, porn, web artifacts or incorrect LID. One must notice that, while this dataset contains texts from the web, it can not be used as a benchmark of noisy web data, because the texts were heavily cleaned and preprocessed before their annotation was performed. We believe however that HPLT-LID may be valuable for the community, because there exist not many LID benchmarks for some languages presented there..

Table 8 presents evaluation of all models (except OpenLID-v2, because the dataset was created with its help) under consideration on HPLT-LID. This is just another proof that there were no regression in performance of OpenLID-v3 compared to OpenLID-v2.

B Additional Data Sources

Table 10 shows the difference between OpenLID-v2 and OpenLID-v3 in terms of training data sources.

Languages outside OpenLID-v2 class inventory.

We discussed adding natural language other class (languages present in glotlid-corpus other than our

Language	# docs	Language	# docs
ast_Latn	139	hrv_Latn	188
bos_Latn	49	ita_Latn	171
cat_Latn	196	jpn_Jpan	87
cmn_Hans	183	nob_Latn	304
ces_Latn	183	pes_Arab	209
deu_Latn	179	por_Latn	111
ell_Grek	613	rus_Cyrl	362
eng_Latn	126	spa_Latn	300
fra_Latn	139	slk_Latn	127
fin_Latn	335	srp_Cyrl	15
glg_Latn	182	yor_Latn	168
hin_Deva	136	total	4,502

Table 7: HPLT 3.0 manual inspection dataset statistics.

Model	FPR	F1	Precision	Recall
OpenLID-v3, 0.5	4e-5	99.84	99.91	99.78
OpenLID-v3	6e-5	99.84	99.85	99.84
GlottLID, 0.5	4e-5	99.51	99.89	99.16
GlottLID	4e-5	99.55	99.89	99.24
ensemble	1e-5	99.53	99.96	99.14

Table 8: Evaluation on HPLT-LID (n=3,873), with and without softmax threshold.

190+) for collecting ‘garbage’ documents and out-of-train-set languages. However, after looking closer into the data, just sampling from about 1,900 those languages and labeling them all as "other" did not seem a good idea. About 1,500 of those language were low-resource ones where only Bibles were available, reported as non-reliable data by GlotLID’s authors¹¹.

After removing non-reliable sources, 366 languages were left. The smallest language in OpenLID-v2 dataset is Yiddish with 923 samples only. Out of 366, we have 150 languages with more samples than that. Some of them were quite large like Low German (nds_Latn) with 117,676 samples (many languages in OpenLID-v2 dataset have less), national like Romansh (roh_Latn) or have more than a million speakers like Chechnyan (che_Cyrl). For the future OpenLID improvement, we recommend extending the model with these languages as separate labels instead of pushing them all into "other".

Not all languages with many samples are expected to be common in crawls: some of them

¹¹<https://github.com/cisnlp/GlotLID/blob/main/sources.md>

Model	FPR	F1	Precision	Recall
<i>FastSpell</i>	n=6,809			
OpenLID-v3, 0.5	2.5e-3	91.08	95.72	89.56
OpenLID-v3	2.9e-3	91.22	95.23	90.32
OpenLID-v2, 0.5	3.9e-3	90.04	94.55	89.07
OpenLID-v2	4.2e-3	90.37	94.05	90.11
GlottLID, 0.5	3.0e-3	91.32	95.19	90.24
GlottLID	3.5e-3	91.39	94.63	91.00
ensemble	1.5e-3	90.92	97.11	87.78
<i>FLORES+</i>	n=155,848			
OpenLID-v3, 0.5	7.6e-5	97.72	97.86	97.71
OpenLID-v3	8.6e-5	97.74	97.70	97.89
OpenLID-v2, 0.5	8.9e-5	98.09	98.62	97.98
OpenLID-v2	1.0e-4	98.11	98.46	98.16
GlottLID, 0.5	3.4e-5	98.34	99.52	97.95
GlottLID	3.8e-5	98.38	99.46	98.07
ensemble	2.6e-5	97.01	98.33	96.5
<i>UDHR</i>	n=10,283			
OpenLID-v3, 0.5	1.5e-4	96.14	97.13	96.39
OpenLID-v3,	1.8e-4	96.21	96.87	96.22
OpenLID-v2, 0.5	2.5e-4	96.27	97.26	96.73
OpenLID-v2	3.0e-4	96.21	96.93	96.91
GlottLID, 0.5	6.7e-5	97.39	98.96	96.76
GlottLID	8.7e-5	97.37	98.78	96.87
ensemble	4.5e-5	95.59	98.07	94.58

Table 9: Multilingual benchmarks. F1, precision and recall in %. ‘0.5’ stands for thresholding softmax scores at 0.5. Ensembling experiments are done with thresholded OpenLID-v3 and GlottLID without thresholding. FLORES+ refers to the devtest split, and FastSpell refers to our version that excludes Nynorsk.

are not alive like Classical Chinese (lzh_Hani) or constructed like Interlingua (ina_Latn). Since OpenLID already featured Esperanto and Latin, we believe other such language are also worth being edited.

Also we found more languages present in both Cyrillic and Latin script, like Serbian: Tatar, for example.

Training data for zxx_Zxxx. und_* are random sequences generated from different scripts.¹² zxx_* are non-linguistic noise in different scripts collected from the web.¹³ (We experimented with adding them as separated classes first; while our validation data had no such class, some languages outside our class inventory were classified as one of this classes; the choice between the two was random according to manual inspection.)

C Bosnian, Croatian, Montenegrin and Serbian

Other benchmarks We are aware of BCS news dataset (SETimes) (Ljubešić and Rupnik, 2022a), comprising roughly 9,000 parallel documents written in each of the languages, from a now-inactive news site that published articles in the languages of South-Eastern Europe. However, SETimes data were present among both OpenLID and GlottLID’s training data, but their splitting into sentences seem to be done in different ways. While we succeeded in deduplicating the test data from GlottLID’s training instances, we scored suspiciously high with OpenLID (more than 90% F1), which was never observed on other benchmarks. This is a clear sign that our deduplication has not worked. Thus, we only report GlottLID results here and leave proper evaluation on SETimes for future work. It is an important benchmark, because it comprises multi-sentence documents; average document length for each combination of gold and assigned label is given in Table 11. Generally, the common wisdom that longer documents lead to better classification holds here as well, mainly because longer documents give more opportunities for clear grammatical markers and discriminators to arise. Short sentences are much more likely to be mislabeled, across all models and datasets. Note that Serbian (srp, bottom row) doesn’t follow the pattern of

¹²https://github.com/cisnlp/GlottLID/blob/main/assets/train/gen_und.py

¹³<https://huggingface.co/datasets/cis-lmu/glottlid-corpus/tree/main/zxx>

Language	Sources	Justification
ast_Latn	commonvoice, leipzigwiki, tatoeba	high confusion with other
bam_Latn	ud, all OpenLID-v2 dyu data	high confusion with dyu_Latn
ben_Beng	commonvoice, GlotStoryBook, leipzigwiki, tatoeba, ud	small in HPLT 3.0
bod_Tibt	GlotStoryBook, tatoeba, the newest Wikipedia	high confusion with dzo_Tibt
cat_Latn	commonvoice, GlotStoryBook, leipzigwiki, tatoeba, ud	some confusion with oci_Latn
dan_Latn	GlotStoryBook, leipzigwiki, tatoeba, ud	high confusion with nob_Latn
dzo_Tibt	the newest Wikipedia	high confusion with bod_Tibt
frp_Latn	HFWikipedia	missed in OpenLID-v2, high confusion with lij_Latn
fuv_Latn	tatoeba	high confusion with other
gug_Latn	commonvoice, leipzigwiki, tatoeba, ud	poor in HPLT 3.0
heb_Hebr	leipzigwiki, tatoeba, ud	minor confusion with underrepresented Yiddish
kan_Knda	GlotStoryBook, leipzigwiki, tatoeba	small in HPLT 3.0
kin_Latn	afriqa, commonvoice, GlotStoryBook, tatoeba	high confusion with run_Latn
lat_Latn	tatoeba, ud, GlotStoryBook, leipzigwiki	missed in OpenLID-v2
lij_Latn	the newest Wikipedia	poor in HPLT 3.0
oci_Latn	removing OpenLID’s ‘pilar’	high confusion with lij_Latn
run_Latn	masakhanews, tatoeba, the newest Wikipedia	high confusion with kin_Latn
rus_Cyrl	commonvoice, GlotStoryBook, leipzigwiki, tatoeba, ud	high confusion with other
som_Latn	GlotStoryBook, leipzigwiki, masakhanews, tatoeba	high confusion with other
srp_Latn	ud, setimes, tatoeba	missed in OpenLID-v2, high confusion with Croatian and Bosnian
ssw_Latn	GlotStoryBook, Vukuzenzele, the newest Wikipedia	high confusion with other
sun_Latn	leipzigwiki, nusa, tatoeba	high confusion with other
tam_Tam1	commonvoice, GlotStoryBook, leipzigwiki, tatoeba, ud	small in HPLT 3.0
zxx_Zxxx	zxx_*, und_*	missed in OpenLID-v2

Table 10: OpenLID-v3 added/removed training data, compared to OpenLID-v2; added data were taken from GlotLID-corpus, if another not specified.

(gold/pred)	bos	hrv	srp
bos	1,922	571	-
hrv	534	1,833	-
srp	6,129	468	3,694

Table 11: GlotLID evaluation on SETimes; average document length by character count, across correctly and incorrectly labeled data.

longer documents being more likely to be correctly labeled. In addition to performance on Latin-script Serbian generally being lower, this discrepancy is likely due to common lexical overlap between Bosnian and Serbian, exacerbated by document length, and the model seemingly giving preference to the bos label, discussed further in this section.

We also acknowledge the existence of the COPA (Rupnik et al., 2023) and Heritage BCS (Romić, 2024) datasets, but, as the Serbian portion of COPA uses the standard Cyrillic script only, it was inappropriate for our evaluation; and Heritage BCS data comprise transcripts of bilingual speakers’ conversations, there is too much second-language interference in the documents for meaningful evaluation on LID task.

Montenegrin GlotLID predictions Interestingly, there are a handful of mislabeled documents

that GlotLID labels as Montenegrin (cnr) or Chakavian (ckm, a dialect of Croatian). Both the first¹⁴ and the second¹⁵ feature nonstandard syntax and repetitive *i* conjunctions and adjectives; both are, respectively, typical of Chakavian and Montenegrin poetry, which feature disproportionately in training data.

HPLT-LID reannotation details 201 samples were predicted by OpenLID-v2 to be Serbian Cyrillic, out of which 151 were pre-annotated to be false positives in the corresponding repository, but turned out to be missing annotation rather than false. 21 samples were detected to be translationese from Serbian Cyrillic and annotated with a new tag tra_Zxxx; 3 samples were found to be only valid in Montenegrin and annotated with cnr_Latn.

D Romance Languages of Italy and France

Table 12 lists the Francoprovençal and Occitan translations of the UDHR and their corresponding URLs.

¹⁴“Dijagnoza je točna, nisam i siguran da je i terapija, nisam siguran da je i terapija ispravna.”

¹⁵“Bijele noći, pa normalno da su bijele noći smetale, sve što je bijelo smetalo je, u mraku se puno bolje rade ovakvi poslovi.”

Dialect	URL
Occitan Auvergnat	http://efele.net/udhr/d/udhr_auv.txt
Occitan Languedocien	http://efele.net/udhr/d/udhr_lnc.txt
Occitan Provençal	http://efele.net/udhr/d/udhr_prv.txt
Francoprovençal Fribourg	http://efele.net/udhr/d/udhr_oci_1.txt
Francoprovençal Savoie	http://efele.net/udhr/d/udhr_oci_2.txt
Francoprovençal Vaud	http://efele.net/udhr/d/udhr_oci_3.txt
Francoprovençal Valais	http://efele.net/udhr/d/udhr_oci_4.txt

Table 12: Occitan and Francoprovençal UDHR translations and their URLs.

E Scandinavian Languages

Table 13 presents ensemble results combining GlotLID (no threshold) and OpenLID-v3 (softmax threshold 0.5). Table 14 shows performance comparison for different versions of OpenLID with GlotLID on the SLIDE, Nordic DSL, FLORES+ devtest, and UDHR test sets with and without thresholding.

F Hierarchical Models

Hierarchical LID systems decompose the classification task into a series of decision steps, typically progressing from broad language groups to the individual languages or variants (Jauhiainen et al., 2019; Agarwal et al., 2023). The motivation for these approaches is that closely related languages are naturally more challenging to discriminate for LID models (Goutte et al., 2014).

We trained hierarchical fastText-based models and evaluated their effect on the following language groups: Scandinavian languages (Danish, Faroese, Icelandic, Norwegian Bokmål, Norwegian Nynorsk, and Swedish), South Slavic languages (Bosnian, Croatian, Serbian in Latin script), along with varieties of Arabic (Egyptian, Levantine, Mesopotamian, Moroccan, Najdi, Standard Arabic, Ta’izzi-Adeni, and Tunisian), and Persian (Dari and Iranian Persian). For each group, we train a specialized language classifier on the subset of the OpenLID-v3 corpus composed of these languages. During inference, every time the base model predicts a language from one of the four groups, we replace the base model prediction with the prediction of the specialized model.

The specialized models follow the same architecture and training procedure as OpenLID-v3. For each language group, we use the corresponding language-specific subset of the OpenLID-v3 training data.

We have not found a significant difference in the performance of OpenLID-v3 and the hierarchical models in any of the related language groups when evaluating on FLORES+ or UDHR. We hypothesize that the capacity of the large LID model is not saturated yet and therefore there is nothing to be gained by using a subset of the training data. Instead, effort should be put into securing larger and higher-quality annotated data for these language groups.

G List of Supported Languages

Table 15 shows the list of language labels supported by OpenLID-v3.

Dataset	Loose Acc.	Exact Acc.	NB F1	DA F1	NN F1	SV F1	Other F1
SLIDE	<i>94.62</i>	<i>92.30</i>	92.52	<i>90.53</i>	91.52	<i>98.21</i>	91.70
Nordic DSL	93.66	93.66	<i>89.41</i>	<i>94.89</i>	97.01	95.54	92.54
FLORES+	99.97	99.97	<i>97.56</i>	<i>99.35</i>	<i>98.45</i>	99.95	99.99
UDHR	99.99	99.99	98.39	99.17	<i>99.13</i>	100.00	100.00

Table 13: Ensembling GlotLID (no threshold) and OpenLID-v3 (softmax threshold 0.5) on Scandinavian languages. Bold indicates improvement over both individual models. Italic indicates improvement over OpenLID-v3 only.

Thres.	Model	Loose Acc.	Exact Acc.	NB F1	DA F1	NN F1	SV F1	Other F1
<i>SLIDE</i>								
No	OpenLID-v3	94.46	90.55	94.33	89.27	93.61	96.30	93.13
	OpenLID-v2	94.39	90.36	94.07	89.60	94.15	95.89	93.11
	OpenLID-v1	93.61	89.77	93.91	87.56	93.03	95.51	91.42
	GlottLID	97.20	93.40	95.83	93.37	94.84	98.23	98.05
0.5	OpenLID-v3	95.55	91.87	94.25	91.60	93.52	97.62	95.16
	OpenLID-v2	95.63	91.83	94.16	91.58	94.12	97.38	95.29
	OpenLID-v1	94.81	91.14	94.00	89.26	93.36	96.84	93.76
	GlottLID	97.09	93.45	95.56	93.40	94.52	98.22	97.73
	n samples	6,950	6,950	2,098	677	1,628	1,250	1,745
<i>Nordic DSL (from 50k split)</i>								
No	OpenLID-v3	94.71	94.71	89.29	93.63	94.24	96.06	97.36
	OpenLID-v2	94.32	94.32	88.63	92.89	93.79	95.67	97.30
	OpenLID-v1	94.17	94.17	88.73	93.64	92.78	95.37	97.15
	GlottLID	96.19	96.19	94.54	95.40	96.84	95.87	97.21
0.5	OpenLID-v3	94.30	94.30	89.17	93.83	94.36	96.11	95.96
	OpenLID-v2	93.96	93.96	88.49	92.84	94.30	95.65	96.07
	OpenLID-v1	93.88	93.88	88.47	93.65	93.16	95.41	96.13
	GlottLID	96.10	96.10	94.48	95.51	96.85	95.96	96.86
	n samples	74,800	74,800	14,960	14,960	14,960	14,960	14,960
<i>FLORES+ devtest</i>								
No	OpenLID-v3	99.97	99.97	97.53	99.11	98.22	99.90	100.00
	OpenLID-v2	99.97	99.97	97.19	98.91	98.08	99.90	100.00
	OpenLID-v1	99.97	99.97	97.63	98.86	98.57	99.75	100.00
	GlottLID	99.98	99.98	98.24	99.60	98.60	100.00	100.00
0.5	OpenLID-v3	99.97	99.97	97.22	99.26	98.32	99.95	100.00
	OpenLID-v2	99.97	99.97	96.84	98.86	98.17	99.95	100.00
	OpenLID-v1	99.97	99.97	97.53	98.86	98.62	99.85	100.00
	GlottLID	99.98	99.98	98.19	99.55	98.60	100.00	100.00
	n samples	216,568	216,568	1,012	1,012	1,012	1,012	212,520
<i>UDHR</i>								
No	OpenLID-v3	99.75	99.75	93.85	97.44	82.19	78.21	99.88
	OpenLID-v2	99.74	99.74	96.12	85.71	95.80	73.49	99.87
	OpenLID-v1	99.79	99.79	83.22	95.31	91.20	88.41	99.90
	GlottLID	99.97	99.97	99.20	96.00	99.15	99.19	99.99
0.5	OpenLID-v3	99.97	99.97	98.39	98.36	97.44	98.39	99.99
	OpenLID-v2	99.97	99.97	98.41	97.56	98.28	97.60	99.99
	OpenLID-v1	99.97	99.97	98.41	99.19	95.00	99.19	99.99
	GlottLID	99.99	99.99	99.20	98.36	100.00	100.00	100.00
	n samples	27,757	27,757	62	61	58	61	27,515

Table 14: Comprehensive comparison of language identification models across all test datasets with and without softmax thresholding. Bold values indicate best performance for each metric within each dataset/threshold combination.

Lang. code	Name	Lang. code	Name	Lang. code	Name	Lang. code	Name
ace_Arab	Achinese	fon_Latn	Fon	lij_Latn	Ligurian	slk_Latn	Slovak
ace_Latn	Achinese	fra_Latn	French	lim_Latn	Limburgan (Limburger, Limburgish)	slv_Latn	Slovenian
afr_Latn	Afrikaans	frp_Latn	Arpitan (Franco-provençal)	lin_Latn	Lingala	smo_Latn	Samoan
als_Latn	Tosk Albanian	fur_Latn	Friulian	lit_Latn	Lithuanian	sna_Latn	Shona
amh_Ethi	Amharic	fuv_Latn	Nigerian Fulfulde	lmo_Latn	Lombard	snd_Arab	Sindhi
ara_Arab	Arabic	gaz_Latn	West Central Oromo	ltg_Latn	Latgalian	som_Latn	Somali
asm_Beng	Assamese	gla_Latn	Scottish Gaelic	ltz_Latn	Luxembourgish (Letzeburgesch)	sot_Latn	Southern Sotho
ast_Latn	Asturian (Asturleonese, Bable, Leonese)	gle_Latn	Irish	lua_Latn	Luba-Lulua	spa_Latn	Spanish (Castilian)
awa_Deva	Awadhi	glg_Latn	Galician	lug_Latn	Ganda	srd_Latn	Sardinian
ayr_Latn	Central Aymara	gug_Latn	Paraguayan Guaraní	luo_Latn	Luo (Kenya and Tanzania) (Dholuo)	srp_Cyrl	Serbian
azb_Arab	South Azerbaijani	guj_Gujr	Gujarati	lus_Latn	Lushai	srp_Latn	Serbian
azj_Latn	North Azerbaijani	hat_Latn	Haitian (Creole, Haitian Creole)	lvs_Latn	Standard Latvian	ssw_Latn	Swati
bak_Cyrl	Bashkir	hau_Latn	Hausa	mag_Deva	Magahi	sun_Latn	Sundanese
bam_Latn	Bambara	heb_Hebr	Hebrew	mai_Deva	Maithili	swe_Latn	Swedish
ban_Latn	Balinese	hin_Deva	Hindi	mal_Mlym	Malayalam	swh_Latn	Swahili (individual language) (Kiswahili)
bel_Cyrl	Belarusian	hne_Deva	Chhattisgarhi	mar_Deva	Marathi	szl_Latn	Silesian
bem_Latn	Bemba (Zambia)	hrv_Latn	Croatian	min_Latn	Minangkabau	tam_Taml	Tamil
ben_Beng	Bengali	hun_Latn	Hungarian	mkd_Cyrl	Macedonian	taq_Latn	Tamasheq
bho_Deva	Bhojpuri	hye_Armn	Armenian	mlt_Latn	Maltese	taq_Tfng	Tamasheq
bjn_Arab	Banjar	ibo_Latn	Igbo	mni_Beng	Manipuri	tat_Cyrl	Tatar
bjn_Latn	Banjar	ilo_Latn	Iloko	mos_Latn	Mossi	tel_Telu	Telugu
bod_Tibt	Tibetan	ind_Latn	Indonesian	mri_Latn	Maori	tgk_Cyrl	Tajik
bos_Latn	Bosnian	isl_Latn	Icelandic	mya_Mymr	Burmese	tha_Thai	Thai
bug_Latn	Buginese	ita_Latn	Italian	nld_Latn	Dutch (Flemish)	tir_Ethi	Tigrinya
bul_Cyrl	Bulgarian	jav_Latn	Javanese	nno_Latn	Norwegian Nynorsk	tpi_Latn	Tok Pisin
cat_Latn	Catalan (Valencian)	jpn_Jpan	Japanese	nob_Latn	Norwegian Bokmål	tsn_Latn	Tswana
ceb_Latn	Cebuano	kab_Latn	Kabyle	npi_Deva	Nepali (individual language)	tso_Latn	Tsonga
ces_Latn	Czech	kac_Latn	Kachin (Jingpho)	nso_Latn	Pedi (Northern Sotho, Sepedi)	tuk_Latn	Turkmen
cjk_Latn	Chokwe	kam_Latn	Kamba (Kenya)	nus_Latn	Nuer	tum_Latn	Tumbuka
ckb_Arab	Central Kurdish	kan_Knda	Kannada	nya_Latn	Chichewa (Chewa, Nyanja)	tur_Latn	Turkish
cmn_Hans	Mandarin Chinese	kas_Arab	Kashmiri	oci_Latn	Occitan (post 1500)	twi_Latn	Twi
cmn_Hant	Mandarin Chinese	kas_Deva	Kashmiri	ory_Orya	Odia (Oriya (individual language))	uig_Arab	Uighur (Uyghur)
crh_Latn	Crimean Tatar (Crimean Turkish)	kat_Geor	Georgian	pag_Latn	Pangasinan	ukr_Cyrl	Ukrainian
cym_Latn	Welsh	kaz_Cyrl	Kazakh	pan_Guru	Panjabi (Punjabi)	umb_Latn	Umbundu
dan_Latn	Danish	kbp_Latn	Kabiyè	pap_Latn	Papiamentu	urd_Arab	Urdu
deu_Latn	German	kea_Latn	Kabuverdianu	pbt_Arab	Southern Pashto	uzn_Latn	Northern Uzbek
dik_Latn	Southwestern Dinka	khk_Cyrl	Halh Mongolian	plt_Latn	Plateau Malagasy	vec_Latn	Venetian
dzo_Tibt	Dzongkha	khm_Khmr	Khmer (Central Khmer)	pol_Latn	Polish	vie_Latn	Vietnamese
ekk_Latn	Standard Estonian	kik_Latn	Kikuyu (Gikuyu)	por_Latn	Portuguese	war_Latn	Waray (Philippines)
ell_Grek	Modern Greek (1453-)	kin_Latn	Kinyarwanda	quy_Latn	Ayacucho Quechua	wol_Latn	Wolof
eng_Latn	English	kir_Cyrl	Kirghiz (Kyrgyz)	ron_Latn	Romanian (Moldavian, Moldovan)	xho_Latn	Xhosa
epo_Latn	Esperanto	kmb_Latn	Kimbundu	run_Latn	Rundi	ydd_Hebr	Eastern Yiddish
eus_Latn	Basque	kmr_Latn	Northern Kurdish	rus_Cyrl	Russian	yor_Latn	Yoruba
ewe_Latn	Ewe	knc_Arab	Central Kanuri	sag_Latn	Sango	yue_Hant	Yue Chinese
fao_Latn	Faroese	knc_Latn	Central Kanuri	san_Deva	Sanskrit	zgh_Tfng	Standard Moroccan Tamazight
fas_Arab	Persian	kor_Hang	Korean	sat_0lck	Santali	zsm_Latn	Standard Malay
fij_Latn	Fijian	ktu_Latn	Kituba (Democratic Republic of Congo)	scn_Latn	Sicilian	zul_Latn	Zulu
fil_Latn	Filipino (Pilipino)	lao_Lao	Lao	shn_Mymr	Shan	zxx_Zxxx	No linguistic content (Not applicable)
fin_Latn	Finnish	lat_Latn	Latin	sin_Sinh	Sinhala (Sinhalese)		

Table 15: OpenLID-v3 supported languages (model class labels).