

Language ID in the Wild: Unexpected Challenges on the Path to a Thousand-Language Web Text Corpus

Isaac Caswell, Theresa Breiner, Daan van Esch, Ankur Bapna
Google Research, 1600 Amphitheatre Parkway, Mountain View, CA 94043
{icaswell, tbreiner, dvanesch, ankurbpn}@google.com

Abstract

Large text corpora are increasingly important for a wide variety of Natural Language Processing (NLP) tasks, and automatic language identification (LangID) is a core technology needed to collect such datasets in a multilingual context. LangID is largely treated as solved in the literature, with models reported that achieve over 90% average F1 on as many as 1,366 languages. We train LangID models on up to 1,629 languages with comparable quality on held-out test sets, but find that human-judged LangID accuracy for web-crawl text corpora created using these models is only around 5% for many lower-resource languages, suggesting a need for more robust evaluation. Further analysis revealed a variety of error modes, arising from domain mismatch, class imbalance, language similarity, and insufficiently expressive models. We propose two classes of techniques to mitigate these errors: wordlist-based tunable-precision filters (for which we release curated lists in about 500 languages) and transformer-based semi-supervised LangID models, which increase median dataset precision from 5.5% to 71.2%. These techniques enable us to create an initial data set covering 100K or more relatively clean sentences in each of 500+ languages, paving the way towards a 1,000-language web text corpus.

1 Introduction

Thousands of languages are spoken in our world (Eberhard et al., 2019), but technologies like machine translation (MT) and automatic speech recognition (ASR) are only available in about 100 of them. As internet access becomes increasingly common with the spread of smartphones (Biggs, 2017), bringing technologies that can help lower language and literacy barriers to more languages is ever more important.

Unfortunately, bringing language technologies to more languages is costly, as for many technologies, extending to an additional language has generally required the use of large parallel labeled datasets. For example, ASR systems are usually trained on large sets of audio recordings and transcriptions, while MT systems have historically needed a set of bilingual sentence pairs. Increasingly, small parallel datasets do exist for many languages (Mayer and Cysouw, 2014; Agić and Vulić, 2019; Artetxe et al., 2020; Ardila et al., 2020), but those resources were either produced at high cost, or are restricted to narrow domains. Parallel resources, which rarely occur naturally, remain scarce for most languages.

Monolingual text data, which is more commonly produced, is also used in building out language technologies: for example, in training language models, which are used in many applications ranging from next-word prediction in keyboard input software (Ouyang et al., 2017) to ASR and MT (Buck et al., 2014). Historically, though, a monolingual text corpus by itself has not been sufficient to build ASR and MT systems in a new language: at least some parallel data was typically necessary.

Recently, however, significant progress has been made in cross-lingual learning for NLP tasks (Klementiev et al., 2012; Ammar et al., 2016; Lample and Conneau, 2019; Pfeiffer et al., 2020): for example, some approaches appear capable of extending machine translation models to new languages with only monolingual data (Artetxe et al., 2017; Lample et al., 2017; Siddhant et al., 2020), and similar findings have been reported for other NLP tasks (Hu et al., 2020). For ASR it is possible to combine a

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

target-language language model with an acoustic model from a phonologically similar language, with no need for parallel datasets of audio recordings and transcriptions (Prasad et al., 2019). Such approaches are likely to get even more effective with nearly-universal acoustic models (Li et al., 2020) and more scalable grapheme-to-phoneme modeling approaches (Deri and Knight, 2016; Mortensen et al., 2018; Bleyan et al., 2019; Ritchie et al., 2019; Ritchie et al., 2020; Lee et al., 2020). Even if more work is needed to establish when such approaches will work well (Marchisio et al., 2020; Artetxe et al., 2020; Wu and Dredze, 2020), having useful monolingual text corpora across languages is clearly a prerequisite to exploring such approaches further. Additionally, using techniques such as LaBSE (Yang and Feng, 2020), parallel corpora can also be constructed from monolingual corpora.

Unfortunately, it has proven challenging to derive highly multilingual text corpora from the web (Artetxe et al., 2020). One commonly cited reason is that most web content is written in widely-spoken languages like English and Mandarin¹. Still, previous work has shown that the web contains labeled and unlabeled data in thousands of languages (Scannell, 2007; Prasad et al., 2018). Since most web pages do not have any language labels attached, previous efforts to build web text corpora often rely at least in part on crawling selected URLs and top-level domains in each language, or use popular n-gram Language Identification (LangID) models like FastText (Grave, 2017) to target a limited number of languages (Goldhahn et al., 2012; Ortiz Suárez et al., 2019). However, previous work (Section 2) has shown that it is possible to build highly accurate LangID systems covering 1,000+ languages.

Thus, aiming to build a 1,000-language web text corpus, we trained a similar large-coverage LangID model, and used it in a large web crawl. However, we found that such LangID systems do not deliver useful results in a real-world web-crawl scenario. To address this, we make the following contributions:

1. We demonstrate that LangID is much less “solved” than frequently believed, and popular n-gram modeling techniques (used for all existing web crawl corpora) have especially serious problems
2. We categorize common problems LangID models fall prey to (Section 3)
3. We present two improvements over existing approaches: tunable-precision wordlist-filtering and Semi-Supervised Transformer models (Section 4)
4. We propose alternative evaluation metrics that better estimate the quality of LangID models from the perspective of web-mining (Section 5) and perform a deep, 600-language web-crawl (Section 6)

This work focuses on monolingual corpora, but the problems described also apply to parallel texts, and it is straightforward to extend the improvements described here to parallel data crawling.

2 LangID Approaches for Web Corpora

To create text corpora in as many languages as possible, we needed a broad-coverage, accurate LangID model for our web crawl. We cover existing work and describe our model, built along similar lines.

2.1 Previous Implementations

A rich literature exists on building text corpora from the web: for example, the *Web as Corpus* workshops have focused on the challenges around identifying relevant pages, extracting clean text, content de-duplication, and many other relevant topics (Barbaresi et al., 2020; Jakubíček et al., 2020). We use an internal web crawler, which is equipped with robust text extraction and de-duplication features, and focus on expanding its LangID component.

A comprehensive recent survey on LangID is Jauhiainen et al. (2018). Naturally, LangID systems have been applied to web crawls before: Buck et al. (2014) published n-gram language models for 175 languages based on Common Crawl data. The Corpora Collection at Leipzig University (Goldhahn et al., 2012) and the Corpus of Global Language Use (Dunn, 2020) offer corpora in 252 and 148 languages. The largest language coverage is probably An Crúbadán, which does not leverage LangID, and found (small amounts of) web data in about 2,000 languages (Scannell, 2007). Our work is probably most

¹We think existing statistics on the distribution of languages on the web should be taken with a grain of salt, as they were likely gathered using highly imperfect language identification models, as discussed in this paper.

similar to OSCAR (Ortiz Suárez et al., 2019) and CCNet (Wenzek et al., 2019), which mined Common Crawl data for 166 and 174 language varieties respectively. However, we believe depth of mining and LangID robustness can limit the quality of datasets produced by these projects: a preliminary inspection of the (often small) low-resource language corpora produced by these LangID-based projects discovers the sort of data noise we describe in this paper, which may render them unusable for NLP applications. These Common-Crawl based datasets are also smaller than our final, filtered dataset, which is $\approx 20x$ larger than CCNet and $\approx 180x$ larger than OSCAR for shared low-resource languages (see Appendix D).

One relevant LangID implementation appearing in the above works is Dunn (2020), achieving an F1 above 0.95 for 464 languages, and offering a thorough evaluation on different data sources and domains. The only LangID systems with higher coverage that we are aware of are those developed by Brown (2012; 2013; 2014), with the most recent version covering as many as 1,366 language varieties, with accuracy above 99%. These numbers are impressive, but as we will see, even such high accuracy on test sets will not suffice to derive useful monolingual corpora from a real-world web crawl.

2.2 Our LangID Implementation

The LangID model we built is similar in approach to previously described systems: we use an n-gram based CLD3 model (Bakalov et al., 2016), consisting of a single hidden layer feed-forward neural network on bag-of-n-gram features and script-count features, which we trained on an aggregation of proprietary and publicly available text corpora, covering 1,629 language varieties, with an average of 800K tokens per language. Some of the data came from sources with language tags like Wikipedia, while another subset was created using a text elicitation task where we prompted native speakers to write sentences in their language (van Esch et al., 2019). For some languages, we also relied on data extracted by Corpus Crawler (Brawer, 2017), a tool which mines text from sites with known in-language content. Using these corpora, we trained several LangID models, on increasingly large sets of languages. As Table 1 demonstrates, even highly multilingual models achieved good F1 scores on held-out test sets.

We balanced the data to have the same size dataset for each language before training. Since the relatively uncommon languages we are targeting have little web data compared to languages like English, balancing the data makes sense in order to have a high-enough recall model to get whatever scarce data there might be on the web for less common languages. Additionally, practically speaking, weighting training data according to the estimated prevalence of each language on the web at large—for example, with orders of magnitude more English examples than Quechua examples—would likely make model training difficult from a computational and stability perspective. However, it is worth stressing that evaluating a model on balanced data overestimates the performance of a model on the highly imbalanced web, especially with respect to precision, as we will see in Section 3.1.

Coverage	Avg. F1	Med. F1
212 lang.	96.1%	98.2 %
1629 lang.	90.4%	97.9%

Table 1: LangID model performance: macro-average F1 and median F1

3 Failure Modes of LangID Models on Web Text

Despite our LangID models performing well on the held-out test sets, when applied on real-life web data, the models were not as accurate as we had expected. We performed an initial limited crawl with a 648-language model, but some quick evaluations showed that the results were highly noisy, so we performed a full crawl on $\approx 100B$ documents with a 224-language model to isolate the problems for closer analysis. This model had comparable performance to the models in Table 1, with median F1 of 96.8 on held-out eval sets. As first-pass filtering, we performed **document-consistency filtering**: we ran the LangID model on every sentence in each document, and then took the most commonly predicted language as the document language. We only kept sentences where the sentence-level and document-level labels matched. All datasets were also de-duplicated. This approach may have decreased recall on multilingual pages, but it reduced the severe noise problems, and helped reduce disk storage needs.

While we expected some accuracy loss due to the domain mismatch between clean training data and



Pred. Language	Mined "Sentence" purporting to be in this language	Noise class
Manipuri		General noise
Twi (Akan)	me: why you hyyin, why you always hyyin	General noise
Varhadi	Òyáèè èè, áóòà- éyòèyòy àèèè àèáí Èàáóá àyèèèè òyìy- yáyá-yòèáú èèy áó íyñè [...]	Misrendered PDF
Aymara	Orilyzewuhubys ukagupixog axiqyh asozasuh uxilitidobyq osoqalelohan [...]	Non-Unicode font
Balinese	As of now  is verified profile on Instagram.	Boilerplate
Cherokee	••••• ALL mY ThØRAs GREW bACK As fLØWERs ••••• SWEET BΛBIES n DØGS	Creative use of Unicode
Oromo	My geology essay introduction essay on men authoring crosswords	Unlucky frequent n-gram
Pular	MEEOW	Repeated n-grams
Chechen	Жи рновский ... Жи рновский районный Ф естив аль Т О Сов	A N T S P E A K
Kashmiri	ਸ਼.	Short/ambiguous
Nigerian Pidgin	This new model features a stronger strap for a secure fit and increased comfort.	High-resource cousin
Uyghur	نۇرسۇلتان نازاربايىۋ قىتايىڭ قازاقستانداغى مەشھۇر	Out-of-model cousin
Dimli	The S<b class="b2">urina<b class="b1">m toa<b class="b3">d is [...]	Deliberately Obfuscated

Table 2: Examples of several representative classes of noise in our initial web-crawl corpora.

noisy web text (Dunn, 2020), even after document-consistency filtering the LangID labels were so noisy that the corpora for the majority of languages in our crawl were unusable for any practical NLP task. Table 2 presents some representative samples of noise. Beyond various kinds of noise, we also found a high number of unexpected misclassifications, as in the Oromo case in Table 2. The following sections detail important classes and sources of noise.

3.1 Massive Class Imbalances: 99% Accuracy Is Not Enough

Precision, unlike recall or false positive rate (FPR)², is a function of the class balance in a dataset. Measuring precision on a balanced dataset may give misleading impressions about real-world performance. For example, consider a LangID model that has 99% precision, 99% recall, and 0.01% FPR on a particular language on a balanced development set. Imagine however that there are 100 billion pages on the web, of which 10,000 are in the target language: in this scenario, the resulting web-crawled dataset will be mostly out-of-language, containing just under *a tenth of a percent* of sentences in the target language (see calculations in Appendix B)—insufficient for most NLP applications. Yet this assumes a relatively low FPR; for languages with a high FPR with respect to a much more common language, like Nigerian Pidgin with English, the situation is even more dire.

As can be seen from this example, calculations of precision (and by extension, F1) are misleading when applied to real-world data with different class balances than the development set. In the general case, for a classifier with recall r and false positive rate f , if we estimate that the language of interest constitutes $x\%$ of the total web text, we get:

$$\text{precision}_{\text{crawl}} = \frac{xr}{xr + (1-x)f} \quad (1)$$

Therefore, any evaluation of LangID models should also report the false positive rate (ideally with respect to major languages on the internet, like English) along with their precision and recall. This class-imbalance effect exacerbates the problems described in the following sections.

3.2 General Internet Noise and Creativity

There are many kinds of web noise that are known to cause problems both with LangID and in downstream tasks, such as abbreviations (“g2g”, “hbu”), leetspeak (“n00b”), hashtags (“#99problems”), or

²In the two-class case. FPR depends on the balance of the other classes with respect to each other, but not on the balance of the target class with respect to all other classes. Per-language FPR (e.g. percent of English sentences classified as Nigerian Pidgin) is truly balance-independent.

non-standard Unicode encodings (like a LATIN CAPITAL LETTER W instead of a CYRILLIC CAPITAL LETTER WE). Some of these problems can be handled automatically (Prasad et al., 2018; Chua et al., 2018). However, our efforts in scaling the LangID models in our web crawl to hundreds of languages uncovered greater depths to internet noise, alongside even more creative ways of using text. As a result of the sheer size of the web, any small pathologies of a LangID model are hugely magnified: we observed that our models tend to pick up on particular genres of internet noise for each separate language, resulting in corpora for some languages that mostly showcase a rich array of particular types of oddities.

For example, in our initial crawls, what purported to be the corpus for Varhadi picked up large amounts of badly-encoded PDFs; Aymara and Turkmen were made up mostly of misrendered non-Unicode text; Dimli had mostly invalid HTML; Dogri offered a rich array of Zalgo-like ornamentation; Fula was awash in URLs; Ilocano caught vast amounts of garbled Javascript; and Zhuang captured German sentences involving the Unicode SOFT HYPHEN character. In each of these cases, sadly the majority of the crawled corpus actually consisted of the class of noise that the LangID classifier decided to assign to these languages—unfortunately drowning out any in-language sentences in the corpora.

In another interesting twist, one might expect that languages which are written in scripts that are not used for any other language would have clean corpora, as the unique connection between the script and the language means that any LangID model gets 100% F1 on development sets. However, this underestimates the creativity of the internet: the Cherokee syllabary, for example, contains characters that look similar to Latin characters, which are consequently repurposed to give words in other languages an aesthetic effect (see example in Table 2), while other scripts, such as Balinese, are used commonly for purely decorative purposes alongside content in entirely unrelated languages. Some script-unique languages like Divehi do yield high-precision corpora right from the get-go, but they are the lucky few.

3.3 Artifacts from Character N-gram Modeling

Many error modes seem to be direct consequences of n-gram count based models, and are also common in public corpora crawled using n-gram models like FastText (Grave, 2017)—Appendix E explores these phenomena in the OSCAR (Ortiz Suárez et al., 2019) corpus. Here are a few important classes of pathologies we discovered; see Table 2 for examples of each, and Appendix C for frequency statistics:

1. **Unlucky overlap of frequent n-grams with high-prevalence languages:** Token frequencies in natural text follow a power law distribution (Zipf, 1935), so that the most common n-grams in a language will be present in a majority of all of its sentences. If one of these common n-grams happens to occur in a sentence in a different language, LangID models can over-trigger. We observed this with Oromo, where 50% of the crawled dataset was actually English sentences containing the word “essay” *at least three times*, misleading the model due to high counts for the n-grams “essa”, “ess”, “sa”, “a”, “e”, “s”, and “y”, all of which are top Oromo n-grams (see Appendix Table 12).
2. **Repeated n-graaaaaaams:** By repeating an n-gram sequence an arbitrary amount, which is rare in clean training text but common on the internet, the class probability of a language may be ramped up, even if the language is clearly wrong—cf. adversarial examples (Goodfellow et al., 2015).
3. **A N T S P E A K :** A surprisingly common internet phenomenon is to find text with space-separated characters, `l i k e t h i s` (Channing, 2020). Standard n-gram models—or even SentencePiece models (Kudo and Richardson, 2018)—can’t handle this without special-casing. This affects about one to two languages per major script: we found that most of our “Chechen” data was actually `R u s s i a n`, most of our “Lambadi” `T e l u g u`, our “Santali” `B e n g a l i`, and some of our “Sepedi” `E n g l i s h`.

3.4 Languages with High-Prevalence Cousins

Languages with High-Prevalence Cousins is a specific, quite common case of the Class Imbalance problem, which requires somewhat different techniques to mitigate (see Section 4). Crawling the web for a low-resource language (“*target language*”) that is closely related to a language that is highly prevalent on the internet (“*distractor language*”) can yield a dataset consisting mostly of the distractor language. A particularly salient example is Nigerian Pidgin (i.e. Naija, ‘pcm’) and English (‘en’), which are similar

enough (see Appendix Table 11 for examples) that typical LangID models will have high false positive rates between the two. Because of the prevalence of English on the internet, along with this high degree of confusability, building a high-precision web-crawled text corpus for languages like Nigerian Pidgin is exceedingly difficult.

3.5 Languages with Out-of-Model Cousins

A variant on the above are languages that are not supported by the LangID model, which interfere with related languages that are supported. For example, a majority of our Uyghur crawl was actually Kazakh and Kyrgyz in the Arabic script; our model had been trained to recognize Kazakh and Kyrgyz, but only in the Cyrillic alphabet. Table 2 gives an example Kazakh sentence that was labeled as Uyghur.

3.6 Unrepresentative Training Data

Sometimes training data may be *too clean* to be accurate on out-of-domain, noisy web data; yet other times it may be too noisy, too homogeneous, or contain systematic biases. For example, for some languages, training data (especially data sourced from Wikipedia) had high quantities of special characters and templated data (esp. from censuses). Templated data may be harmful for n-gram models, by skewing the token distributions away from that of normal text, though there is some evidence that neural models may be less affected by token distributions than by latent structure (Papadimitriou and Jurafsky, 2020). Other training data may also have issues; for instance, in our elicited Chechen data, the CYRILLIC LETTER PALOCHKA (not found on many keyboards) was represented with the ASCII digit “1”. Our model therefore may not handle Chechen text containing the correct code point, or other substitutes, very well.

4 Improving LangID Precision on Web Text

Monolingual web-text corpora afflicted by the issues described in Section 3 will likely prove unusable for practical purposes. We report on two distinct approaches we found helpful in improving precision.

4.1 Tunable-precision Filtering with Curated Wordlists

We experimented with token-based filtering techniques, which are simple to implement and fast to perform on large corpora. Since the LangID models in our crawl operated on character n-grams, token-based approaches may have complementary behavior and can side-step particular failure modes. For instance, since a sentence with the word “essay” likely contains mostly non-Oromo words, the havoc caused by the n-gram “essa” described in Section 3.3 is neatly sidestepped by checking against a curated list of known Oromo words. Such filtering approaches have the added benefit of *tunable precision*, allowing us to adjust the cleanliness of our corpora depending on the noise tolerance of downstream tasks.

Percent-Threshold filtering

The simplest approach to token-based filtering is to remove any sentence where less than $x\%$ of its tokens appear in a clean list of known words for the language, such as one would find in a standard dictionary. We used in-house lists with a median of $\approx 15\text{K}$ words per language, which were obtained through frequency sorting followed by human curation. The one parameter for filtering—the percentage of in-vocabulary words—provides a simple, interpretable way to tune for precision/recall. We call this method **Percent-Threshold Wordlist Filtering**.

TF-IDF based filtering

Percent-Threshold Wordlist Filtering is effective for a majority of the problems we saw, where the text is nonsense or in an entirely different language, but it will *not* help where the mislabeled text is in a similar language, as in Nigerian Pidgin (‘pcm’), which has very high lexical overlap with English (‘en’)—meaning that such filtering will still retain most English sentences, and fail to increase precision. This problem will occur with any language that has high lexical overlap with a major language. Where there is extensive borrowing of loanwords, the languages may even be unrelated, as for Chuvash and Russian.

Some words, however, are highly effective language markers: for example, “wetin” is common in Nigerian Pidgin, but does not occur in English. We therefore propose to keep any sentence that has at

least one word from a small list of common tokens that are *distinctive* to that particular language, and are not shared with its more prevalent cousins. We call this **Disjunctive Wordlist Filtering**.

First, we perform TF-IDF, where each “document” is our LangID training set. However, this suffers one crucial flaw: the idf formulation of TF-IDF weights each document equally, so a word will be equally penalized if it occurs in English or in K’iche’. For practical purposes, we care mainly about filtering out common distractor-language text on the internet, so we only want to penalize those languages.

This motivates a simple variant on TF-IDF which we call TF-IIF, or **Term Frequency-Inverse Internet Frequency**. This measure is the ratio of the frequency of a token in our per-language corpus (TF) with the frequency of that token across the entire internet (IIF), which we approximate from a sample of 7 million randomly selected web sentences. In practice we find that performance improves slightly when accounting for both IDF and IIF, yielding the TF-IDF-IIF score. Formally, for a token t in a language l , with a frequency function $f(\text{term}, \text{corpus})$ and language-specific corpora D_l :

$$\text{tf-idf-iif}_{t,l} = \text{tf}_{t,l} * \text{idf}_{t,l} * \text{iif}_t = f(t, D_l) \log \left(\frac{1}{\sum_{l' \neq l} 1\{t \in D_{l'}\}} \right) \frac{1}{f(t, \text{internet})} \quad (2)$$

With a ranked TF-IDF-IIF list for each language, we then pick the top N words for each language such that we have at least $r\%$ recall on our dev sets. While it is tempting to choose the same r for all languages (e.g. 95%), different languages can behave quite differently with such filters, with small changes in recall sometimes leading to large changes in precision. We had best results by choosing $r \in [0.75, 1.0]$, and then determining the ideal precision-recall trade-off on a per-language basis. With this paper, we publicly release TF-IDF-IIF wordlists we used, covering the top 100 tokens for each of about 500 languages³.

4.2 Semi-Supervised LangID

A separate approach from filtering is to improve our original LangID model. Utilizing large unsupervised text corpora to improve the quality of neural networks has become increasingly important in NLP (Devlin et al., 2018; Wang et al., 2018). Following this line of work, we use the noisy data crawled with our n-gram LangID model to improve the quality of our LangID system by leveraging self-supervised approaches, yielding a Semi-Supervised LangID system (**SS-LID**).

Specifically, following the text-to-text self-supervised approach outlined in Raffel et al. (2019), we train a Transformer Big model (Vaswani et al., 2017) by sampling equally from the crawled data from 212 languages. We co-train this self-supervised task with the LangID task in a text-to-text setting, with the hope of improving the quality of LangID on noisy open-domain web text. To reduce the confounding effect of using a higher capacity transformer, we train a baseline transformer on just the LangID task.

We evaluate these SS-LID models and compare against the n-gram based LangID model in Table 3. In addition to

model	F1	prec.	rec.	FPR
NG-LID ₂₁₂	96.05	94.93	97.64	0.01079
XF-LID ₂₁₂	97.51	97.26	97.82	0.00849
$\Delta\epsilon$	36.9%	46.0%	7.4%	21.4%
SS-LID ₂₁₂	98.03	97.61	98.55	0.00683
$\Delta\epsilon$	50.2%	52.9%	38.4%	36.7%
SS-LID ₆₂₄	97.52	97.86	97.45	0.00610
$\Delta\epsilon$	37.3%	57.8%	-8.3%	43.5%

Table 3: Performance of n-gram LangID model, Transformer LangID model (XF-LID) and Semi-supervised models (SS-LID) trained on either 212 or 624 languages. Scores are averaged over the shared 212 languages.

F1, precision, and recall, we report FPR, whose importance we discussed in Section 3.1. All values are macro-averaged over the shared 212 languages. To distinguish between apparently well-performing models we also report the relative error reduction with respect to the n-gram model, which for an error metric ϵ we define as $\Delta\epsilon = \frac{\epsilon_b - \epsilon_t}{\epsilon_b}$, where ϵ_b is the baseline model error and ϵ_t the test model error.

We see that the Transformer LangID model outperforms the n-gram model by a large margin, especially on precision and FPR. The SS-LID models improve further upon this model, notably with a 40%

³<https://github.com/google-research-datasets/TF-IDF-IIF-top100-wordlists>

Language	Unfiltered		Threshold		Disjunctive		SS-LID ₆₂₄	
	<i>p</i>	<i>r</i>	<i>p</i>	<i>r</i>	<i>p</i>	<i>r</i>	<i>p</i>	<i>r</i>
Aymara	2.1	100	86.2	98.3	76.4	92.9	94.6	99.3
Bhojpuri	4.0	100	3.0	100	4.0	93.0	83.0	98.5
Chechen	24.0	100	84.0	99.9	49.0	91.8	98.0	99.9
Cherokee	16.0	100	95.0	100	97.0	90.6	47.0	100
Chuvash	5.0	100	3.0	43.2	22.0	93.5	5.0	99.5
Divehi	98.8	100	98.6	99.0	99.1	91.4	98.9	99.6
Guarani	4.0	100	12.0	99.0	44.0	92.1	23.0	98.8
Oromo	5.0	100	78.0	98.0	80.0	91.8	33.0	99.5
Surjapuri	31.3	100	45.9	97.1	61.0	88.2	60.3	95.6
Swiss German	2.0	100	2.0	98.7	2.0	92.1	16.0	95.5
Tamazight	6.0	100	42.0	98.8	35.0	91.3	42.0	99.8
Twi (Akan)	49.0	100	83.0	100	83.0	92.4	82.0	93.5
Zhuang	1.0	100	59.0	85.9	3.0	92.6	15.0	98.3
Median	5.5	100	52.5	98.5	47.5	92.0	71.2	98.7

Table 4: Comparison of filtering approaches for a few languages: percent-threshold filtering ($x = 20\%$), disjunctive TF-IDF-IIF filtering ($r = 90\%$), and filtering with a Semi-supervised LangID model. We report 1. human-judged LangID precision over the crawl (percent of in-language sentences), and 2. recall of this method on our held-out eval sets. Best precision is bolded. Full table in Appendix.

reduction in FPR. It is worth noting that these improvements are on the **clean** eval set, despite the additional training objective being on the noisy web crawl. We suspect the improvements are even greater on web-type data, which is partially validated by the evaluation on web-text in Section 5.

5 Evaluating LangID Filtering Methods on Web-Text

5.1 Evaluation Methodology: Principles and Suggestions

Ideally, LangID models would be evaluated on a large, noisy test set, representative of real-life web data. Since such sets do not currently exist, we recommend having human annotators evaluate crawled corpora to ensure quality meets the threshold for downstream use (which will vary per application). For automatic metrics, we suggest focusing on false positive rate and recall rather than precision and recall, and comparing models using relative error reduction to amplify differences between apparently highly-performant models, as we did above in Section 4.2.

5.2 Evaluating our Systems

We asked human annotators to evaluate LangID quality for our web-crawled text in a subset of the languages. First, we filtered the web crawl with several methods. We then randomly sampled 100-1,000 sentences from each of these filtered data sets, and asked annotators (who were fluent speakers, or who spoke a closely related language) to indicate whether each sentence was in the target language.

Table 4 presents the results of this evaluation for a selection of languages (full results on seventeen languages in Appendix Table 5). For each language, we show the precision of the method from the human annotations, and the recall of the same filter on our clean dev sets. For the percent-threshold filtering we evaluated a threshold of 20%, and for the disjunctive wordlist filtering we used the top N TF-IDF-IIF words per language such that the recall on our held-out eval set was at least 90%.

We see that the initial datasets were extremely noisy, with a median value of 5% of sentences being in-language. The filtering methods drastically increased the percentage of correctly LangID’d sentences, with values of up to 99% in-language, while maintaining high recall. However, the best filtering method varies widely by language. The neural SS-LID model has the highest precision for Bhojpuri and Swiss German, both of which also suffer most from the High-Prevalence-Cousin issue among these languages.

However, it does much more poorly than wordlist-based approaches on Oromo and Cherokee. In the latter case, we found that SS-LID was unable to discard English sentences written in Cherokee syllabics.

It is worth re-emphasizing that the thresholds in Table 4 were chosen somewhat arbitrarily for the purpose of illustration. Since precision is tunable in the word-based approaches, precision can be increased further, though at growing cost to recall—a trade-off to make depending on downstream noise tolerance.

For Guinea-Bissau Creole, which has both a High-Prevalence Cousin (Portuguese) and an Out-of-Model Cousin (Papiamentu), none of our filtering methods were effective (see Appendix). Swiss German, in the same situation, barely scraped by. Future work should investigate additional techniques for such cases—although the most effective solution may be as simple as using a hand-curated TF-IDF-IRF list, which looked promising in preliminary experiments in Nigerian Pidgin.

6 Web-crawled Dataset and Comparison with other Public Datasets

Using the above methods⁴, we performed a deep crawl of the web (touching >100B webpages) with a 600-language LangID model. Using percent-threshold filtering⁵ we made a recall-focused dataset, then post-filter with a SS-LID model for high precision, yielding a larger, cleaner set than is found in similar corpora. More details and comparisons to public corpora (OSCAR, CCNet) are in Appendices E and D.

7 Future Work

Our approach yielded usable monolingual text corpora in ≈ 600 languages. Internal user experience research suggests the web may now contain at least some amount of monolingual text in thousands of languages, so we plan to scale up with more multilingual LangID models, like our 1,629-language model.

Truly covering the linguistic richness of the web will also need crawling approaches to be fine-tuned further. Text for some languages may only be found in PDF files (Bustamante et al., 2020), and some scripts are commonly represented in non-Unicode fonts—such as Kruti Dev for Devanagari, requiring separate detection for conversion into Unicode-encoded Devanagari (Singh and Goyal, 2013). Applying OCR may also help handle non-Unicode text, and can uncover textual content within images. And many languages that are not officially written in the Latin alphabet have informal transliterated orthographies (Roark et al., 2020); our models can identify the most common ones, but we could cover more.

Finally, our work focused on a web crawl, but many new internet users primarily use their language online on social media platforms and in chat messages (Soria, 2018; van Esch et al., 2019). Other work has looked at applying LangID to social media (Jaech et al., 2016; Blodgett et al., 2017; Vo and Khoury, 2019). Our techniques should help improve LangID accuracy in this challenging domain, too.

8 Conclusion

Language Identification (LangID) is by no means a solved problem, and n-gram models are much worse than popularly believed. We trained LangID models covering up to 1,629 languages, but found that even seemingly high-quality models (> 95 F1) were nearly unusable in practice for low-resource languages. We described and analyzed several major issues encountered in applying LangID to a real-life web crawl. These practical problems included large amounts of noise, much of which appears to be natural language and can't be easily filtered out; insufficient expressiveness of n-gram models; issues with related languages; and a massive class imbalance problem, meaning that even 99% F1 can be insufficient.

To solve these issues, we developed two major improvements to our LangID system: tunable-precision filtering methods (for which we release wordlists in about 500 languages) and semi-supervised neural models. These allowed us to create usable monolingual text corpora across hundreds of languages based on our deep web crawl, with much more and cleaner data per language than previously published approaches. Such corpora hold great promise for bringing technologies like MT and ASR to more languages, and we believe it should be possible to use the approaches we outlined to create monolingual corpora in many more languages, which should help extend language technology even further.

⁴Our process is also summarized in Appendix K for those interested in replicating.

⁵In this case, we used larger wordlists than those used for the analysis above, in order to stress recall.

Acknowledgements

We would like to thank Diana Akrong, Alex Rudnick, Mikhail Donolin, Maxim Krikun, Hakim Sidahmed, and Landis Baker for help with human evaluations of the LangID models, as well as Vera Axelrod, Jason Riesa, and Wolfgang Macherey for useful advice and reviews. We also want to specifically thank Onome Ofoman for her consultation and advice about Nigerian Pidgin.

References

- Željko Agić and Ivan Vulić. 2019. JW300: A wide-coverage parallel corpus for low-resource languages.
- Waleed Ammar, George Mulcaire, Yulia Tsvetkov, Guillaume Lample, Chris Dyer, and Noah A. Smith. 2016. Massively Multilingual Word Embeddings. *ArXiv*, abs/1602.01925.
- Rosana Ardila, Megan Branson, Kelly Davis, Michael Kohler, Josh Meyer, Michael Henretty, Reuben Morais, Lindsay Saunders, Francis Tyers, and Gregor Weber. 2020. Common voice: A massively-multilingual speech corpus. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 4218–4222, Marseille, France, May. European Language Resources Association.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George F. Foster, Colin Cherry, Wolfgang Macherey, Zhifeng Chen, and Yonghui Wu. 2019. Massively multilingual neural machine translation in the wild: Findings and challenges. *CoRR*, abs/1907.05019.
- Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2017. Unsupervised neural machine translation. *arXiv preprint arXiv:1710.11041*.
- Mikel Artetxe, Sebastian Ruder, Dani Yogatama, Gorka Labaka, and Eneko Agirre. 2020. A Call for More Rigor in Unsupervised Cross-lingual Learning. In *ACL*.
- Anton Bakalov, Alex Salcianu, Andy Golding, Chris Alberti, Daniel Andor, David Weiss, Emily Pitler, Greg Coppola, Jason Riesa, Kuzman Ganchev, Michael Ringgaard, Nan Hua, Ryan McDonald, Slav Petrov, Stefan Istrate, and Terry Koo. 2016. Compact Language Detector v3 (CLD3), October.
- Adrien Barbaresi, Felix Bildhauer, Roland Schäfer, and Egon Stemle, editors. 2020. *Proceedings of the 12th Web as Corpus Workshop*, Marseille, France, May. European Language Resources Association.
- P. Biggs. 2017. The state of broadband: Broadband catalyzing sustainable development. Technical report, International Telecommunications Union and Broadband Commission for Sustainable Development and UNESCO.
- Harry Bleyan, Sandy Ritchie, Jonas Fromseier Mortensen, and Daan van Esch. 2019. Developing Pronunciation Models in New Languages Faster by Exploiting Common Grapheme-to-Phoneme Correspondences Across Languages. In *Proceedings of Interspeech 2019*.
- Su Lin Blodgett, Johnny Tian-Zheng Wei, and Brendan O’Connor. 2017. A Dataset and Classifier for Recognizing Social Media English. *ACL*.
- Sascha Brawer. 2017. Corpus Crawler. <https://github.com/google/corpuscrawler>, September.
- Ralf D Brown. 2012. Finding and identifying text in 900+ languages. *Digital Investigation*, 9:S34–S43.
- Ralf D Brown. 2013. Selecting and weighting n-grams to identify 1100 languages. In *International Conference on Text, Speech and Dialogue*, pages 475–483. Springer.
- Ralf Brown. 2014. Non-linear mapping for improved identification of 1300+ languages. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 627–632, Doha, Qatar, October. Association for Computational Linguistics.
- Christian Buck, Kenneth Heafield, and Bas Van Ooyen. 2014. N-gram Counts and Language Models from the Common Crawl. In *LREC*, volume 2, page 4. Citeseer.
- Gina Bustamante, Arturo Oncevay, and Roberto Zariquiey. 2020. No Data to Crawl? Monolingual Corpus Creation from PDF Files of Truly low-Resource Languages in Peru. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 2914–2923, Marseille, France, May. European Language Resources Association.

- Cornelia Channing. 2020. A Biologist Explains Why She and 1.8 Million Others Are Role-Playing as Ants on Facebook. <https://slate.com/culture/2020/05/facebook-ants-roleplay-coronavirus-biologist-interview.html>. Accessed: 2020-06-01.
- Mason Chua, Daan van Esch, Noah Coccaro, Eunjoon Cho, Sujeet Bhandari, and Libin Jia. 2018. Text Normalization Infrastructure that Scales to Hundreds of Language Varieties. In *Proc. of the 11th edition of the Language Resources and Evaluation Conference, LREC 2018, 7-12 May 2018, Miyazaki, Japan*.
- Aliya Deri and Kevin Knight. 2016. Grapheme-to-phoneme models for (almost) any language. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 399–408, Berlin, Germany, August. Association for Computational Linguistics.
- 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*.
- Jonathan Dunn. 2020. Mapping languages: the Corpus of Global Language Use. *Language Resources and Evaluation*, April.
- David M. Eberhard, Gary F. Simons, and Charles D. Fennig, editors. 2019. *Ethnologue: Languages of the World*. SIL International, Dallas, Texas, 22 edition.
- Dirk Goldhahn, Thomas Eckart, and Uwe Quasthoff. 2012. Building large monolingual dictionaries at the leipzig corpora collection: From 100 to 200 languages. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 759–765, Istanbul, Turkey, May. European Language Resources Association (ELRA).
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and Harnessing Adversarial Examples. *ICLR*.
- E. Grave. 2017. Language identification, October. Accessed: October 22, 2020.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization. *arXiv preprint arXiv:2003.11080*.
- Aaron Jaech, George Mulcaire, Shobhit Hathi, Mari Ostendorf, and Noah A. Smith. 2016. Hierarchical Character-Word Models for Language Identification.
- Miloš Jakubiček, Vojtěch Kovář, Pavel Rychlý, and Vit Suchomel. 2020. Current challenges in web corpus building. In *Proceedings of the 12th Web as Corpus Workshop*, pages 1–4, Marseille, France, May. European Language Resources Association.
- Tommi Jauhiainen, Marco Lui, Marcos Zampieri, Timothy Baldwin, and Krister Lindén. 2018. Automatic Language Identification in Texts: A Survey. *CoRR*, abs/1804.08186.
- Alexandre Klementiev, Ivan Titov, and Binod Bhattarai. 2012. Inducing Crosslingual Distributed Representations of Words. In *COLING*.
- Taku Kudo and John Richardson. 2018. SentencePiece: A simple and language independent subword tokenizer and detokenizer for Neural Text Processing. *arXiv preprint arXiv:1808.06226*.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *arXiv preprint arXiv:1901.07291*.
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2017. Unsupervised machine translation using monolingual corpora only. *arXiv preprint arXiv:1711.00043*.
- Jackson L. Lee, Lucas F.E. Ashby, M. Elizabeth Garza, Yeonju Lee-Sikka, Sean Miller, Alan Wong, Arya D. McCarthy, and Kyle Gorman. 2020. Massively multilingual pronunciation modeling with WikiPron. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 4223–4228, Marseille, France, May. European Language Resources Association.
- Xinjian Li, Siddharth Dalmia, Juncheng Li, Matthew Lee, Patrick Littell, Jiali Yao, Antonios Anastasopoulos, David R. Mortensen, Graham Neubig, Alan W Black, and Florian Metze. 2020. Universal Phone Recognition with a Multilingual Allophone System.
- Kelly Marchisio, Kevin Duh, and Philipp Koehn. 2020. When Does Unsupervised Machine Translation Work?

- Thomas Mayer and Michael Cysouw. 2014. Creating a Massively Parallel Bible Corpus. In *Ninth International Conference on Language Resources and Evaluation, LREC 2014, 26-31 May 2014, Reykjavik, Iceland*, pages 3158–3163. Research Unit Quantitative Language Comparison, Philipps University of Marburg.
- David R. Mortensen, Siddharth Dalmia, and Patrick Littell. 2018. EpiTran: Precision G2P for many languages. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May. European Language Resources Association (ELRA).
- 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)*, Cardiff, United Kingdom, July.
- T. Ouyang, D. Rybach, F. Beaufays, and M. Riley. 2017. Mobile Keyboard Input Decoding with Finite-State Transducers. *CoRR abs/1704.03987*.
- Isabel Papadimitriou and Dan Jurafsky. 2020. Pretraining on Non-linguistic Structure as a Tool for Analyzing Learning Bias in Language Models. *arXiv preprint arXiv:2004.14601*.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020. MAD-X: An Adapter-based Framework for Multi-task Cross-lingual Transfer.
- Manasa Prasad, Theresa Breiner, and Daan van Esch. 2018. Mining training data for language modeling across the world’s languages. In *Proc. The 6th Intl. Workshop on Spoken Language Technologies for Under-Resourced Languages*, pages 61–65.
- Manasa Prasad, Daan van Esch, Sandy Ritchie, and Jonas Fromseier Mortensen. 2019. Building Large-Vocabulary ASR Systems for Languages Without Any Audio Training Data. In *Proceedings of Interspeech 2019*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Sandy Ritchie, Richard Sproat, Kyle Gorman, Daan van Esch, Christian Schallhart, Nikos Bampounis, Benoit Brard, Jonas Fromseier Mortensen, Millie Holt, and Eoin Mahon. 2019. Unified Verbalization for Speech Recognition & Synthesis Across Languages. In *Proceedings of Interspeech 2019*.
- Sandy Ritchie, Eoin Mahon, Kim Anne Heiligenstein, Nikos Bampounis, Daan van Esch, Christian Schallhart, Jonas Fromseier Mortensen, and Benoit Brard. 2020. Data-Driven Parametric Text Normalization: Rapidly Scaling Finite-State Transduction Verbalizers to New Languages. In *Proceedings of the 1st Joint SLTU and CCURL Workshop (SLTU-CCURL 2020)*, page 218–225, Language Resources and Evaluation Conference (LREC 2020), Marseille.
- Brian Roark, Lawrence Wolf-Sonkin, Christo Kirov, Sabrina J. Mielke, Cibu Johny, Isin Demirsahin, and Keith Hall. 2020. Processing South Asian languages written in the Latin script: the dakshina dataset. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 2413–2423, Marseille, France, May. European Language Resources Association.
- K. P. Scannell. 2007. The Crúbadán Project: Corpus building for under-resourced languages. In *3rd Web as Corpus Workshop, 2007, Louvain-la-Neuve, Belgium*.
- Aditya Siddhant, Ankur Bapna, Yuan Cao, Orhan Firat, Mia Chen, Sneha Kudugunta, Naveen Arivazhagan, and Yonghui Wu. 2020. Leveraging monolingual data with self-supervision for multilingual neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2827–2835, Online, July. Association for Computational Linguistics.
- UmrinderPal Singh and Vishal Goyal. 2013. Font identifier and Unicode converter for Hindi. *Nepalese Linguistics*.
- Claudia Soria. 2018. Digital Language Survival Kit. <http://www.dldp.eu/en/content/digital-language-survival-kit>.
- Daan van Esch, Elnaz Sarbar, Tamar Lucassen, Jeremy O’Brien, Theresa Breiner, Manasa Prasad, Evan Elizabeth Crew, Chieu Nguyen, and Françoise Beaufays. 2019. Writing Across the World’s Languages: Deep Internationalization for Gboard, the Google Keyboard. Technical report.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.

- Duy Tin Vo and Richard Khoury. 2019. Language Identification on Massive Datasets of Short Message using an Attention Mechanism CNN.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzman, Armand Joulin, and Edouard Grave. 2019. CCNet: Extracting High Quality Monolingual Datasets from Web Crawl Data, 11.
- 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, July. Association for Computational Linguistics.
- Yinfei Yang and Fanxiaoyu Feng. 2020. Language-Agnostic BERT Sentence Embedding, August. Accessed: September 25, 2020.
- George Zipf. 1935. *The Psychology of Language*. Houghton-Mifflin.

A Complete human evaluation results

A more complete version of Table 4 is given here in Table 5, containing the full set of seventeen languages we evaluated. The only additional information it shows over Table 4 is the percentage of the web-crawl each method filters out, for more context into how these methods will behave in practice. (Keep in mind that, while the precision and % filtered rows are measured on the noisy web crawl, the recall is measured on the held-out eval set.)

Language	Unfiltered			Threshold			Disjunctive			SS-LID		
	<i>p</i>	<i>r</i>	<i>flt.</i>	<i>p</i>	<i>r</i>	<i>flt.</i>	<i>p</i>	<i>r</i>	<i>flt.</i>	<i>p</i>	<i>r</i>	<i>flt.</i>
Ahirani*	49.1	100	0.0	38.2	100	27.2	46.0	90.7	49.8	96.4	98.6	86.6
Aymara	2.1	100	0.0	86.2	98.3	98.3	76.4	92.9	98.2	94.6	99.3	98.1
Bashkir*	33.1	100	0.0	84.9	95.2	62.1	91.5	91.9	67.7	89.7	99.4	61.0
Bhojpuri	4.0	100	0.0	3.0	100	28.3	4.0	93.0	57.6	83.0	98.5	97.0
Chechen	24.0	100	0.0	84.0	99.9	73.9	49.0	91.8	69.1	98.0	99.9	78.7
Cherokee	16.0	100	0.0	95.0	100	86.9	97.0	90.6	87.7	47.0	100	68.6
Chuvash	5.0	100	0.0	3.0	43.2	56.3	22.0	93.5	89.5	5.0	99.5	54.8
Divehi	98.8	100	0.0	98.6	99.0	2.7	99.1	91.4	27.2	98.9	99.6	0.0
Guarani	4.0	100	0.0	12.0	99.0	77.4	44.0	92.1	93.5	23.0	98.8	85.4
G.B. Creole* †	0.0	100	0.0	0.0	100	18.5	0.0	93.6	36.1	0.0	92.9	76.3
Kinyarwanda*	37.3	100	0.0	79.6	93.0	58.1	88.1	91.9	62.0	90.9	98.8	60.8
Oromo	5.0	100	0.0	78.0	98.0	99.0	80.0	91.8	99.0	33.0	99.5	97.4
Surjapuri	31.3	100	0.0	45.9	97.1	34.7	61.0	88.2	51.8	60.3	95.6	77.9
Swiss German	2.0	100	0.0	2.0	98.7	70.6	2.0	92.1	43.3	16.0	95.5	88.5
Tamazight	6.0	100	0.0	42.0	98.8	88.2	35.0	91.3	78.6	42.0	99.8	88.5
Twi (Akan)	49.0	100	0.0	83.0	100	50.0	83.0	92.4	42.5	82.0	93.5	38.9
Zhuang	1.0	100	0.0	59.0	85.9	98.9	3.0	92.6	84.6	15.0	98.3	88.2
Median	5.5	100	0.0	52.5	98.5	60.1	47.5	92.0	64.9	71.2	98.7	82.6

Table 5: More complete comparison of different filtering approaches for different languages. For each example language, we report 1. the precision of the crawl (percent of in-language sentences), as judged by human raters over a sample of 100 sentences per filtering method, 2. the recall of this method on our held-out eval sets, and 3. the percentage of the crawl removed by this filtering method. * Starred languages were omitted from the table in the main paper. † G.B. = Guinea-Bissau

B Massive Class Imbalance: Worked Example

This section shows the methodology for the example in Section 3.1, where we examine by way of example a LangID model with 99% precision, 99% recall, and 0.01% FPR for a given language. If we approximate that there are 100 billion pages on the web, of which 10,000 are in a language we are seeking, we can analyze the precision of the web crawl using the quantities of True Positives (TP), True Negatives (TN), False Negatives (FN), and False Positives (FP). For the dataset resulting from the web crawl, we can therefore say that $TN + FP \approx 100B - 100k \approx 100B$, and $TP + FN \approx 100k$. One can now calculate p_{crawl} , the precision on the resulting crawl of the web:

$$TP = \frac{TP}{TP + FN}(TP + FN) = r * (TP + FN) = 0.99 * 10k = 9.9k \quad (3)$$

$$FP = \frac{FP}{TN + FP}(TN + FP) = fpr * (TN + FP) = 0.0001 * 100B = 10M \quad (4)$$

$$p_{crawl} = \frac{TP}{TP + FP} = \frac{9.9k}{9.9k + 10M} \approx 0.1\% \quad (5)$$

C Statistics on languages most affected by different types of noise

Many of the types of noise mentioned in Section 3.2 are hard to quantify without significant extra work. For instance, it would require building special classifiers for misrendered PDFs, non-Unicode fonts, creative use of Unicode, and so on—and it may need a stronger classifier than an n-gram classifier, since after all these are mistakes of an n-gram classifier. Issues like out-of model-cousins are even trickier, probably requiring human ratings. However, some types of noise can be quantified using approximations like the following:

- **A N T S P E A K**: regex match with `/[^] [^] [^] [^] [^] /`
- **n-graaaaams**: regex match with `/((.)\2\2\2\2)/` up to `/((.....)\2\2\2\2)/`
- **HTML**: regex match with `<[a-z/]*>/`
- **http**: regex match with `/http/`
- **Title Case**: `> 5` successive tokens such that `x[0].isupper()` and `x[1:].islower()`
- **essay**: (special for Oromo) regex match with `[Ee]ssay/`
- **misrendered PDF**: contains bigrams along the lines of `{áí, íè, ñò}` etc. or `{^j, j^, ^J}` etc. (basically, we created a very simple bigram classifier on known misrendered PDFs)

Language (Script)	Phenomenon	Percent of crawl
Lambadi (Telu.)	A N T S P E A K	72.1%
Santali (Beng.)	A N T S P E A K	58.2%
Bodo (Beng.)	n-graaaaams	50.9%
Pular	n-graaaaams	26.3%
Avar	HTML	64.2%
Dimli	HTML	93.1%
Fula	http	44.5%
Magahi	http	23.6%
Nigerian Fulfulde	Title Case	64.5%
Balinese	Title Case	63.1%
Oromo	essay	64.4%
Varhadi	misrendered PDF	90.8%
Yucateco	misrendered PDF	74.7%

Table 6: Quantification of the incidence of a few noise phenomena, along with their most affected languages in our web-crawl.

D Details on the web-mined datasets

As described in Section 6, the dataset we mined has two versions, one focused on recall (called `recall` in the table), and one focusing on precision (called `sslid(recall)` in the table). Table 7 compares these two datasets with public benchmarks.

Since the purpose of this crawl was to focus on low-resource languages, we mined a smaller portion of the internet for the ~ 100 highest-resource languages, and did not do any filtering on these languages. For this reason, in addition to the stats on the entire dataset, we report the stats on the dataset omitting the highest-resource 100 languages, to give a fairer approximation of the size of datasets for truly low-resource languages. We also report stats on the languages among those that are shared between the three datasets, again omitting the ~ 100 highest resource languages.

Please note that these datasets are hard to compare to public benchmarks, as they crawl a wider swath of the internet, and are much more highly multilingual. Therefore, the comparison with public data

sources in this table should not be interpreted as giving information about the nature of the filtering methods described in this paper.

metric subset	N Languages			N Sentences			Median Dataset size		
	all	100+	shared	all	100+	shared	all	100+	shared
recall	600	500	59	36B	3200M	3200M	2100K	970K	12000K
SS-LID(recall)	600	500	59	2.8B	600M	740M	200K	100K	1400K
CCNet	174	74	59	70B	4.4M	22M	930K	6K	78K
OSCAR	166	66	59	20B	0.5M	5.4M	200K	1K	8K

Table 7: Comparison between the two versions of our dataset and the public datasets CCNet and OSCAR. Although the statistics look similar on the full dataset, we see that the public datasets are heavily skewed towards higher-resource languages. When excluding the 100 highest-resource languages (“100+”), or looking only at shared low-resource languages (“shared”), we see that the public datasets have 20x to 200x less data than our crawl was able to identify.

E Comparison with OSCAR Corpus

While the analyses in the main paper focused on evaluating the quality of the data we crawled, publicly available datasets have similar issues. This section briefly analyzes the OSCAR corpus (Ortiz Suárez et al., 2019), which, although an excellent resource for many languages, has lower-quality content for some languages. All analyses are performed on the deduplicated OSCAR corpus, which is cleaner.

Please note that it is hard to compare OSCAR directly with our dataset. One notable confound is that the two datasets are drawing from different portions of the web. Another confound is the degree of multilinguality and the subset of languages chosen (this paper tends to focus on longer-tail languages than OSCAR). A further large confound is that OSCAR uses the FastText LangID model (Grave, 2017), which does not upsample training data, and therefore will tend to have lower recall and higher precision.

Applying the heuristic analyses from Section C, we see that repeated ngram and A N T S P E A K issues are also very common in the OSCAR corpus (the other phenomena from Table 6, however, were mostly absent). Table 8 reports the three most affected languages per phenomenon, and Figure 1 shows a representative sample of two of these corpora. In both these cases, the dataset consisted only of such noise, and had no in-language content.

To further analyze the cleanliness of the OSCAR corpus, we performed a similar analysis as in Section 5, to determine the percentage of each dataset that was in-language. Table 9 summarizes these findings, along with the percentage of the corpus remaining after percent-threshold filtering with our wordlists. We only look at the thirty lowest-resource languages in the corpus. We find that the percent in-language varies widely by language, ranging from 0% to 100%. However, many of the corpora have relatively high precision, with the average precision being just over 89%. At the same time, this accords with a low average recall, with the median dataset size being only 37 sentences. It is interesting to note that wordlist-filtering corresponds quite well with human-judged precision, with Pearson’s R of 87.3%.

Language	Phenomenon	% of crawl
Central Bicol	A N T S P E A K	100.0%
Neapolitan	A N T S P E A K	100.0%
Em.-Romagnol	A N T S P E A K	55.8%
Somali	n-graaaaams	88.1%
Cantonese	n-graaaaams	57.1%
Asturian	n-graaaaams	53.0%

Table 8: Most-affected languages in the OSCAR corpus for two common error modes of n-gram models

F Notes on Curated Wordlist Approaches

For languages written in unsegmented scripts (where spaces are not used in between words; for example, Mandarin), leveraging the curated wordlists during the filtering techniques is not as straightforward. When given a sentence to check for valid words, we would first need to run a segmentation model in order to split the sentence into words, but segmentation models need to be trained on specific languages and do not usually support lower-resource languages. To handle languages written in such writing systems, we included all valid characters in the language as part of the wordlist, so that we could fall back to character-level checks for any sentences written in these scripts. This means that any somewhat reasonable language data using the same script will be kept, even if it is a different language.

G Wordlist-based Language ID

For languages with little or no training sentence-level data, even an n-gram LangID model is not practical to train. We therefore additionally explored pure wordlist-based models: specifically, we experimented with a Word-Based LangID system (**WB-LID**), which assigns a LangID label to the sentence by simply counting how many known words appear in the sentence for each possible language and predicting the language with the highest counts, with extra weight granted to “unique words” that appear only in a single language’s wordlist. The simple architecture of WB-LID does not compare to an n-gram LangID model for most languages (Table 10), and we decided not to pursue using the outputs of WBLID as a filter in this work, but this approach seems stable and scalable to more languages, and may be worth exploring in the future as a LangID system for languages where no sentence data can be found to train an n-gram model.

LangID system Comparisons (median F1)	493 Languages	590 Languages
n-gram LangID	97%	96%
Word-Based LangID	75%	76%

Table 10: Performance of the n-gram LangID system vs Word-Based LangID system on development sets. For the dev sets shown in this comparison, we only include languages for which we had both sentence data to train the n-gram model and known wordlists to train the WBLID. We remove any known words from our WBLID system that do not appear in the sentence data used to train the n-gram model. The n-gram model is trained on all sentence data for the supported languages.

H Illustration of the High-Prevalence-Cousin problem

Although the issue of highly similar varieties is very common and may be familiar to speakers of most languages in the world, English-speaking researchers may be less familiar with it, since close relatives of English do not generally receive a lot of attention in the literature. As an illustration, Table 11 gives some examples of Nigerian Pidgin and the English translations. It is clear that a simple classifier might have trouble distinguishing them, especially for more technical sentences.

Nigerian Pidgin	English
abeg, you fit help me?	please, can you help me?
no dey buy wetin we no need	don’t go buying what we don’t need
He don accuse her family say dem inflate di value	He has accused her family of inflating the value
Born August 28, 1991	Born August 28, 1991
Structured, goal-oriented education	Structured, goal-oriented education

Table 11: Examples of Nigerian Pidgin versus English. It is very hard to mine datasets of Nigerian Pidgin from the web, because it is close enough to English that Language ID models and frequent-wordlist filtering methods will pick up a lot of English. In the informal register, like the first few examples, they are more distinguishable, but in the formal, written register they can appear identical.

I Oromo: A Case Study in Unfortunate N-gram Overlap

As alluded to in Section 3.3, Oromo has the peculiar error mode that our n-gram model massively over-triggers with English, despite the two languages bearing little to no resemblance to each other, as a result of the frequent 4-gram “essa”. Table 12 illustrates this further, showing the most common n-grams in true Oromo, in natural English, and in the web-crawl that claimed to be Oromo.

Oromo LID	idx.	“Oromo” crawl	idx.	English LID	idx.
atti	0	ssay	>1000	your	>1000
anii	1	essa	4	have	>1000
akka	2	tion	>1000	with	>1000
eess	3	writ	>1000	what	>1000
essa	4	atio	>1000	here	>1000
jedh	5	mple	>1000	ther	>1000
isaa	6	ment	>1000	tion	>1000
oota	7	ampl	>1000	want	930
kees	8	tive	>1000	like	>1000
itti	9	ting	>1000	thin	>1000

Table 12: Top 10 most common 4-grams in a) Oromo LangID training data, b) the “Oromo” crawl of the web, and c) English LangID training data. Each 4-gram is presented with its index among the top 1,000 most common Oromo 4-grams. We can understand from the n-gram list that the “Oromo” crawl is majority English, overtriggering because of the 4-gram “essa”, from the English word “essay”. In fact, 50% of sentences in the “Oromo” crawl contain the word “essay” at least three times! The other common n-grams in this Table from the “Oromo” crawl are epiphenomenal, reflecting only English words that tend to occur in English sentences about essays.

J Correlation of filtering precision with relevant variables

When do some filtering methods work better than others? We do not have enough data points to make strong statements (N=17), but there are some trends that may be worth commenting on here. In Table 13, we look at the correlation of the precision of unfiltered data and the three proposed filtering methods, and how they correlate with 1) the size of the crawled dataset, and 2) the dialectical relatedness to common languages online. We hypothesize that variable (1) is a combination of variable (2) with non-linguistic noise artifacts, so looking at these two variables can give us an idea of which methods are better at general noise filtering (from train-data pathologies, etc.) and distinguishing related languages.

Unfortunately the “dialectical relatedness to common languages online” is hard to quantify. As a rough approximation, we introduce four heuristic “confusability classes”:

1. **Class 1:** No obviously confusable languages
2. **Class 2:** Confusable low-resource languages or slightly confusable high-resource language
3. **Class 3:** Medium-confusable high-resource language
4. **Class 4:** Very confusable high-resource language

To perform the regression we assign these classes to the values {1, 2, 3, 4}. Per-language assignments are given in Table 14.

Based on the numbers in Table 13, it looks like both wordlist filtering methods perform similarly, and the SS-LID method is noticeably better when languages are more confusable, and possibly slightly worse when there are larger datasets (signalling more confusion with non-linguistic or out-of-domain noise).

	log(n. segs)	confusion rank
unfiltered	-0.71	-0.16
threshold	-0.26	-0.75
disjunctive	-0.21	-0.46
SS-LID	-0.55	0.15

Table 13: Pearson correlation of the precision of three filtering methods (and unfiltered data) with two relevant variables. Number of segments (i.e. number of sentences in the “unfiltered” dataset) is passed through a log transform first, since the size of the unfiltered datasets follows a log distribution. For an explanation of the “confusion rank”, please see the appendix section J and Table 14.

Language	confusion class	Notes / relevant related languages
Divehi	class #1	script unique
Zhuang	class #1	pretty unique orthography
Cherokee	class #2	some confusion with “Cherokee English” etc.
Guarani	class #2	some lexical overlap with Spanish
Tamazight	class #2	some lexical overlap with ar-Latn and Tamasheq
Twi(Akan)	class #2	some lexical overlap with Ewe, Ga, etc.
Kinyarwanda	class #2	high lexical overlap with Rundi etc.
Aymara	class #2	some lexical overlap with Spanish
Oromo	class #2	some lexical overlap with Gedeo, Hamar, Somali etc.
Bashkir	class #3	medium lexical overlap with Russian
Ahirani	class #3	medium lexical overlap with Hindi
Chechen	class #3	medium lexical overlap with Russian
Surjapuri	class #3	medium lexical overlap with Hindi
Chuvash	class #4	high lexical overlap with Russian*
Bhojpuri	class #4	high lexical overlap with Hindi
Guinea-Bissau Creole	class #4	high lexical overlap with Portuguese
Swiss German	class #4	high lexical overlap with German

Table 14: Heuristic judgement of “confusability” for use in the regression in Table 13. Please note that this is not a rigorous quantification of these languages and may contain mistakes. For explanations of the “classes”, please see text. * Note that Chuvash is considered “high” overlap because of polluted training data.

K Complete Recipe

This section is simply a concise description of the steps we took to create our dataset, in the form of suggestions for someone interested in creating a similar dataset.

1. Train LangID model

- (a) **Balance the data first** in order to have higher recall. The distribution of languages in training data may not be representative of the distribution of languages on the web. Temperature sampling (Arivazhagan et al., 2019) may also be a good alternative, in order to decrease over-triggering somewhat.
- (b) If it is computationally feasible to apply a more complex model at inference time, a Transformer-based LangID model (especially co-trained with a self-supervised objective on in-domain text) will have better performance, even if the held-out scores seem only slightly better.
- (c) **Evaluate cannily:** use out-of-domain held-out sets if possible, and pay special attention to the relative reduction in false-positive rate. A model with FPR of 0.1 is much different than one

with FPR of 0.01—don't give up once you reach 95% F1.

2. **Curate wordlists.** If the publicly released wordlists don't suit one's purposes, one could take e.g. the 200 most frequent tokens from the train set, removing words that are also in highly-prevalent languages if desired, like English, Portuguese, Spanish, Russian, German, Chinese, and Hindi. One can skip this step if the Transformer LangID model is good enough, but it will still be useful for tuning the precision of the final datasets, and will still improve for several languages (e.g. in our situation, it was necessary to catch English written in Cherokee script).
3. **Perform the web crawl.** Document-consistency filtering is highly recommended (only output sentences whose sentence-level ID matches the majority sentence-level ID on the page).
4. **Deduplicate the web-crawled data** and filter with wordlists to reach a desired precision.
5. **Look at samples of every language in the dataset!** Even quickly eyeballing the dataset can reveal serious problems. Also consider quickly checking that all the language codes are plausible: for instance, is the `als` data a mix of Tosk Albanian (ISO639-3 code `als`) and Swiss German (which Wikipedia stores under the code `als`)? Or are there some macrolanguage codes in the dataset that cover a superset of other already-covered languages, like Norwegian Bokmal `nb`, Norwegian Nynorsk `nn`, and the macrolanguage code Norwegian `no`?