
Multilingual Spoken Words Corpus

Mark Mazumder Harvard University	Sharad Chitlangia Harvard University	Colby Banbury Harvard University	Yiping Kang University of Michigan	
Juan Ciro Factored.ai/MLCommons	Keith Achorn Intel	Daniel Galvez NVIDIA	Mark Sabini Landing AI	Peter Mattson Google
David Kanter MLCommons	Greg Diamos Landing AI	Pete Warden Google	Josh Meyer Coqui	Vijay Janapa Reddi Harvard University

Abstract

1 Multilingual Spoken Words Corpus is a large and growing audio dataset of spoken
2 words in 50 languages collectively spoken by over 5 billion people, for academic
3 research and commercial applications in keyword spotting and spoken term search,
4 licensed under CC-BY 4.0. The dataset contains more than 340,000 keywords,
5 totaling 23.4 million 1-second spoken examples (over 6,000 hours). The dataset
6 has many use cases, ranging from voice-enabled consumer devices to call center
7 automation. We generate this dataset by applying forced alignment on crowd-
8 sourced sentence-level audio to produce per-word timing estimates for extraction.
9 All alignments are included in the dataset. We provide a detailed analysis of the
10 contents of the data and contribute methods for detecting potential outliers. We
11 report baseline accuracy metrics on keyword spotting models trained from our
12 dataset compared to models trained on a manually-recorded keyword dataset. We
13 conclude with our plans for dataset maintenance, updates, and open-sourced code.

14 1 Introduction

15 Keyword spotting (KWS) is a core technology for consumer-facing voice-enabled interfaces on
16 endpoint devices and is ubiquitous on smartphones (e.g., Siri and Google Voice Assistant). Keyword
17 spotters are typically used as lightweight wake-word engines, constantly running on-device listening
18 for a single phrase. Upon detection, the device may invoke a larger speech model for voice-based
19 tasks such as turning on a light bulb. In recent years, neural network-based KWS has shown great
20 promise [34, 5, 19, 7, 6]. State of the art models have demonstrated high classification accuracy [5],
21 sufficient for wide-scale deployment of KWS across a range of consumer use cases and consumer-
22 facing devices. However, a dearth of open-source, multilingual, and diverse datasets precludes
23 extending the benefits of KWS across the world, especially for under-represented groups.

24 Traditionally, datasets for keyword spotting models require significant manual effort to collect and
25 validate thousands of utterances for each keyword of interest. Consequently, much of the literature
26 relies on existing keyword datasets such as Google Speech Commands [32] and Hey Snips [9]. These
27 datasets are usually monolingual and contain only a handful of keywords in controlled low-noise
28 environments. Moreover, KWS must be robust to a wide range of speaker characteristics (e.g. accents,
29 genders, tone, pitch and stress) and environmental settings. This presents an additional level of
30 challenge in sourcing a corpus that is rich and diverse.

Table 1: Multilingual Spoken Words Corpus (MSWC) compared to existing keyword datasets.

Dataset	Languages	Keywords	Hours	Commercial Use	Keyword Generation
MSWC (This Work)	50	344,286	6601.39	Yes	Automatic
Hey Snapdragon [15]	1	4	1.19	Yes	Manual
Hey Snips [9]	1	1	127	No	Manual
Speech Commands [32]	1	35	27.92	Yes	Manual
Common Voice Single Target [2]	34	14	141	Yes	Manual

31 In this paper, we automate the generation of a multilingual keyword dataset to ensure reproducibility
 32 and maintainability for both academic research and commercial use. We develop the Multilingual
 33 Spoken Words Corpus (MSWC) dataset by applying forced alignment on Common Voice [2], a
 34 crowdsourced speech corpus. Prior work has used forced alignment to automate the extraction of
 35 keywords from speech corpora [14], but challenges in validating the quality of the alignments have
 36 precluded the establishment of a reference corpus. In addition to our dataset, we provide an outlier
 37 metric for each extracted keyword that reflects the quality of the sample. Our contributions include:

- 38 • We release a large multilingual keyword dataset containing over 23.4 million 1-second
 39 spoken examples for over 340,000 keywords from approximately 115,000 source speakers
 40 across 50 different languages, which correspond to over 5 billion speakers across the
 41 planet [11].
- 42 • We provide forced alignments containing per-word timings for each Common Voice source
 43 sentence across all 50 languages.
- 44 • To facilitate comparative benchmarks, we split the corpus into `train`, `dev`, and `test` splits
 45 for each keyword with an 80:10:10 division.
- 46 • We identify anomalous samples via a self-supervised nearest-neighbors strategy. An outlier
 47 metric enables a user to select between a larger noisier dataset, and a smaller cleaner dataset.
- 48 • We quantify the domain gap between manually recorded and automatically extracted key-
 49 words. We establish the relationship between our outlier metric and the ability for a KWS
 50 model trained on recorded keywords to correctly recognize their corresponding extractions.

51 2 Related Work

52 Most keyword spotting datasets tend to be small and often only provide a single language (Table 1).
 53 The Speech Commands dataset [32] is the current standard for keyword spotting research. The dataset
 54 contains 105,829 one-second utterances. However, it supports only 35 words in English. Other
 55 keyword datasets like the Qualcomm Hey Snapdragon Keyword Dataset [15] and the Hey Snips
 56 Dataset [9] target specific “wake words,” limiting potential applications. The Common Voice Single
 57 Target dataset [2] is the only keyword dataset to support a wide number of languages, but it only
 58 contains 14 keywords (digits 0-9 and 4 predefined keywords) per language. Speech synthesis has
 59 been investigated for KWS but can only generate training data in high resource languages [33, 17].

60 The majority of speech datasets focus on sentence length alignment for speech-to-text applications.
 61 Large scale, multilingual datasets like Common Voice [2] and Multilingual LibriSpeech [24] provide
 62 tens of thousands of hours of speech data in multiple languages. However, keyword spotting
 63 applications require individual words for training, and therefore cannot readily take advantage of
 64 these large speech corpora. Forced alignment techniques have been used to generate keyword spotting
 65 training data from larger sentence length datasets [14, 3, 19]. However, [14, 3] only use English
 66 alignments, and [19] only analyzes 440 keywords in 22 languages, whereas our dataset contains
 67 over 340,000 keywords in 50 languages. Additionally, we provide a metric to assess the quality of
 68 extractions (Sec. 4.6), and we plan to regularly update and grow our dataset (Sec. 7).

69 3 Ontology of Multilingual Spoken Words Corpus

70 In this section, we introduce our Multilingual Spoken Words Corpus (MSWC) dataset, which contains
 71 spoken words from 50 different languages, ranging from high-resource languages such as English
 72 and Spanish to low-resource languages such as Oriya (an Indo-Aryan language spoken in the Indian




Figure 1: The number of audio clips and unique keywords for each of the 50 languages in MSWC.

73 state of Odisha) and Divehi (spoken in the Maldives). In total, the MSWC dataset has 23.4 million
 74 unique audio clips (Figure 1). The entirety of the MSWC dataset is open-sourced (including the
 75 forced alignments used to generate our dataset) along with our code.¹ We also provide a detailed
 76 analysis of the MSWC dataset by summarizing properties of the dataset by its constituent languages
 77 and keywords.

78 3.1 Wide-range of Languages

79 Our dataset’s primary goal is to provide keyword spotting training and testing data for languages
 80 with previously limited or non-existent publicly available data, e.g., Italian and Ukrainian with 68
 81 million and 33 million speakers in total respectively. Our dataset contains keyword audio clips for
 82 50 languages. We define the resource level of a language in our dataset by the total number of
 83 hours of audio present. Languages with under 10 hours of spoken word samples are referred to
 84 as *low-resource*, those between 10 and 100 hours are *medium-resource* and languages with over
 85 100 hours are *high-resource*. The MSWC dataset contains spoken word audio for 26 low-resource
 86 languages, 12 medium-resource languages and 12 high-resource languages (Table 2). Out of these 50
 87 languages, the MSWC dataset is, to the best of our knowledge, the only open-source dataset with
 88 spoken word data for 46 languages. Download sizes for each subset are also shown in Table 2.

89 3.2 Unique Keywords and Audio Clips

90 Figure 1 shows the number of unique keywords and audio clips in the MSWC dataset. Our audio
 91 clips are opus-compressed using a single channel with a 48KHz sample rate. For high-resource
 92 languages, as defined in Table 2, the MSWC dataset has on average 23,408 unique keywords and
 93 1,847,608 audio clips per language. To provide a very coarse estimate of the potential outreach of
 94 this subset of our data, we note that these languages are collectively estimated to be spoken by 2.75
 95 billion people worldwide according to Ethnologue [11]. We expect the large swath of keywords and
 96 audio samples in MSWC will facilitate KWS models and capabilities for many domains and use
 97 cases in these languages. Similarly, for medium-resource languages, the MSWC dataset provides on
 98 average 4,054 unique keywords and 109,579 audio clips per language. Several languages, such as
 99 Ukrainian (33M people [11]), are getting their first publicly available spoken word dataset.

100 For low-resource languages, the MSWC dataset contains on average 552 unique keywords and
 101 10,431 audio clips per language. The languages in this low-resource category have the fewest unique
 102 keywords and audio clips due to the limited data currently available in Common Voice. Still, we
 103 expect this data to enable experiments and exploration into keyword-related tasks in these languages
 104 for the first time. Recent work shows that just three to five keyword training examples are sufficient
 105 to fine-tune an embedding model for keyword spotting and achieve a high keyword classification
 106 accuracy across a wide variety of different languages [14, 19]. Therefore, just a handful of keyword
 107 examples can enable unprecedented research and applications for low-resource languages.

¹https://github.com/harvard-edge/multilingual_kws. Corpus is currently available to reviewers.

Table 2: 50 Languages in the MSWC dataset and the size of their corresponding sub-dataset and the number of hours of audio available in MSWC. The languages are organized into 3 resource levels.

Availability	Languages (Size in MSWC, Hours of Audio in MSWC)
Low Resource <10 hours	Arabic (0.1G, 7.6h), Assamese (0.9M, 0.1h), Breton (69M, 5.6h), Chuvash (28M, 2.1h), Chinese (zh-CN) (42M, 3.1h), Divehi (0.7M, 0.04h), Frisian (0.1G, 9.6h), Georgian (20M, 1.4h), Guarani (0.7M, 1.3h), Greek (84M, 6.7h), Hakha Chin (26M, 0.1h), Hausa (90M, 1.0h), Interlingua (58M, 4.0h), Irish (38M, 3.2h), Latvian (51M, 4.2h), Lithuanian (21M, 0.46h), Maltese (88M, 7.3h), Oriya (0.7M, 0.1h), Romanian (59M, 4.5h), Sakha (42M, 3.3h), Slovenian (43M, 3.0h), Slovak (31M, 1.9h), Sursilvan (61M, 4.8h), Tamil (8.8M, 0.6h), Vallader (14M, 1.2h), Vietnamese (1.2M, 0.1h)
Medium Resource >10 & <100 hours	Czech (0.3G, 24h), Dutch (0.8G, 70h), Estonian (0.2G, 19h), Esperanto (1.3G, 77h), Indonesian (0.1G, 11h), Kyrgyz (0.1G, 12h), Mongolian (0.1G, 12h), Portuguese (0.7G, 58h), Swedish (0.1G, 12h), Tatar (4G, 30h), Turkish (1.3G, 29h), Ukrainian (0.2G, 18h)
High Resource >100 hours	Basque (1.7G, 118h), Catalan (8.7G, 615h), English (26G, 1957h), French (9.3G, 754h), German (14G, 1083h), Italian (2.2G, 155h), Kinyarwanda (6.1G, 422h), Persian (4.5G, 327h), Polish (1.8G, 130h), Russian (2.1G, 137h), Spanish (4.9G, 349h), Welsh (4.5G, 108h)




Figure 2: Distribution of word lengths across all 50 languages.

108 3.3 Keyword Length Variation

109 We characterize the relative frequency of word lengths in MSWC by language. Figure 2 shows
 110 a distribution of the number of extractions versus the number of characters per keyword for each
 111 language. We apply a minimum character length of three as part of the data extraction pipeline as a
 112 coarse stop-word filter (Sec. 4.4). We note that violin plot smoothing artifacts appear under three
 113 characters for some languages. The exception is Chinese where we also extract two character words,
 114 as many words in Chinese are two characters; we use the zh-CN language variant. The visualization
 115 is normalized within each language and does not show relative distributions across languages.

116 Figure 2 shows that the keywords in our dataset have a high degree of diversity across all 50 languages.
 117 This diversity is more prominent in medium and high-resource languages. For low-resource languages,
 118 our dataset contains keywords with the number of characters ranging from three to 19, where there
 119 are 619 unique extractions for words with more than 13 characters. Our medium resource language
 120 sub-dataset features keywords that range from three to 29 characters long. There are 253 extractions
 121 for keywords with >19 characters. Our high-resource subset features keywords that range from three
 122 to 34 characters long and 1298 extractions of the extracted keywords are longer than 23 characters.

123 3.4 Keyword Characterization

124 We characterize the parts-of-speech and semantic meanings of keywords in English, Spanish and
 125 Arabic, as representative candidates from high-resource and low-resource languages in Table 2.

126 **Parts-of-Speech** We leverage the “log-linear part-of-speech tagger” [16] in the multilingual
 127 Stanza [25] library to assign part-of-speech (POS) tags to keywords in English, Spanish and Arabic.
 128 Table 3 shows POS tags for keywords in these languages with the number of corresponding clips and
 129 keywords, along with 3 examples (an extended version is provided in the Appendix). We observe
 130 the most common categories are nouns, verbs and adjectives, demonstrating our applicability to

Table 3: Parts of Speech Tags for words in English, Spanish and Arabic. Number of clips (# C) and keywords (# K) in that language. Rows are sorted by the number of keywords in English.

POSTag	# C	# K	English	# C	# K	Spanish	# C	# K	Arabic
Noun	2M	19K	boy, time, man	491K	13K	parte, ciudad, años	2K	236	قرايسدا، باتاكلا، عيش
Proper noun	382K	10K	i'm, i'll, i've	38K	725	del, juan, york	-	-	-
Verb	1M	6K	are, have, had	151k	5k	gran, primera, mismo	2K	102	تناك، تنك، ناك
Adjective	628K	4K	other, more, many	179K	6K	gran, primera, mismo	643	51	ليلقلا، ليمج، بريثكلا
Adverb	683K	1K	there, also, about	56K	378	más, muy, además	74	2	طقف، ك، لذك
Pronoun	644K	97	you, that, his	121k	117	los, las, una	44	5	متيوار، لاق، ي، لهتيوار
Interjection	40K	69	like, please, well	2k	31	hey, joder, adiós	-	-	-
Auxiliary	305K	46	was, can, will	54k	98	fue, son, está	-	-	-
Numeral	121K	35	one, two, three	29k	42	dos, tres, cuatro	217	6	رشع، ثلاث، دحاو

Table 4: Semantic keyword characterization in English, Spanish and Arabic via zero-shot multilingual NLI with representative samples. Number of clips (#C) and keywords (#K). Sorted by #K in English.

Category	# C	# K	English	# C	# K	Spanish	# C	# K	Arabic
Event	3M	1K	School, Preakness	153K	238	Campeonato, Episodio	2K	30	لصح، لاق، ي
Human activity	1M	820	Shooting, Prefers	78K	211	Expanding, Visitando	1K	28	يرتشملا، داغتس
Location	391K	379	Home, County	213K	324	Zona, Pueblo	333	14	لصح، لاق، ي
Name	327K	303	Margot, Cooney	30K	128	Eduardo, Peter	1K	24	يقديعاس، نادا
Animal	131K	281	Sheep, Camel	99K	137	Águila, Especies	1K	25	ة حافت، هييف
Number	301K	279	Second, Six	31K	75	Multiplicar, Primero	493	17	دحاو، قرشع
City	106K	212	York, London	34K	94	Madrid, Berlín	49	4	وساكيو، نايلناطيدع
Technology	52K	191	Television, Videotape	21K	106	Automotriz, Lego	92	5	تاراييسدا، زافلتلا
Culture	52K	180	Popularize, Music	15K	86	Tradicional, Concertista	57	4	قراهمب، يينيبيصدا
Game	48K	131	Play, Kirby's	8K	46	Deportes, Partido	-	-	-
History	52K	119	Stories, Ancient	12K	73	Emperador, Históricas	257	8	ركذتعا، لبيق

131 keyword spotting tasks, as keywords of interest in voice interfaces are often nouns and verbs. Arabic,
 132 a low-resource language, contains 236 unique nouns and 102 unique verbs each with more than
 133 2000 extractions. We provide a significant amount of novel keyword data in low-resource languages
 134 which we expect to enable new advancements. We also conduct POS tag analyses for six additional
 135 languages (German, Greek, Russian, Turkish, Vietnamese, and Chinese) in Appendix Fig. 4, and we
 136 open-source our code to enable users to similarly explore the other languages in MSWC.

137 **Semantic Categorization** We inspect keywords in our dataset by semantic class. We apply a Natural
 138 Language Inference (NLI) model based on RoBERTa [8] on keywords in English, Spanish and Arabic
 139 and categorize each word into predefined categories based on common Wikipedia categories (e.g.,
 140 science, philosophy, religion). Table 4 samples several categories with 2 examples each (expanded
 141 further in Appendix Table 10). We observe 60% of keywords belong to *event*, *human activity*, *name*
 142 and *animal* categories, and find a relatively even distribution among remaining categories. Appendix
 143 Fig. 5 includes six additional languages (German, Greek, Russian, Turkish, Vietnamese, and Chinese),
 144 and as above, we provide code to facilitate user analysis of other languages in our dataset.

145 3.5 Speaker Analysis and Background Noise

146 Our corpus contains an upper bound of 115,000 speakers across 50 languages. The number of
 147 speakers for a given language broadly corresponds with the language’s resource level. The median
 148 speaker count for high, medium, and low-resource languages are 4023, 258, and 54, respectively.
 149 Note that the speakers originally contributed their audio samples to Mozilla Common Voice, subject
 150 to its collection and validation practices. Speakers had discretion in whether their demographic info
 151 was collected during a recording session. If a speaker was not logged in across multiple sessions,
 152 they may appear as multiple speakers. In terms of gender diversity, the MSWC corpus contains 60%
 153 male, 15% female and 25% unknown speakers.

154 Common Voice’s crowd-sourced recordings are performed in uncontrolled settings on computers
 155 and mobile devices, and exhibit background noise (e.g., static, wind, mouse clicks, background
 156 voices) which can improve the robustness of speech models in real-world settings. We note that for
 157 additional robustness, our dataset can be further augmented with sources of synthetic noise [32] or
 158 environmental samples such as babbling, traffic, and crowds [26].

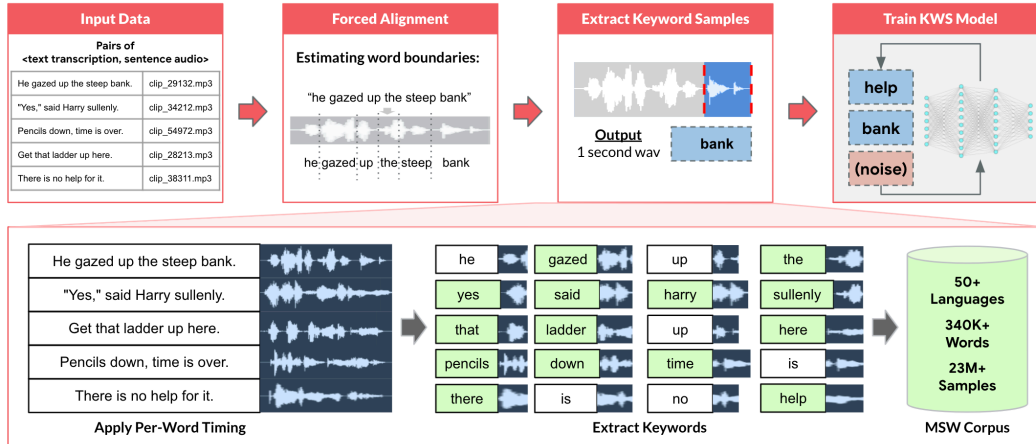


Figure 3: Our automated data extraction pipeline is designed to support multiple languages within the same general workflow, so that we can readily scale to include more than 50 languages over time.

159 4 Constructing the Multilingual Spoken Words Corpus

160 4.1 Data Extraction Pipeline

161 One of our key goals is to develop an automatic and scalable data extraction pipeline so that we
 162 can scale to include more languages and keywords per language over time. To this end, we provide
 163 an overview of our automated keyword extraction pipeline and its constituent elements, which we
 164 used to generate MSWC. Our pipeline is summarized in Fig. 3. The input to our pipeline is one or
 165 more speech datasets consisting of a tuple of (sentence audio, text transcription) pairs, and as output,
 166 we produce: (a) a set of word-aligned timing estimates for the dataset, (b) a dataset of extracted
 167 keywords, and (c) data splits for training, validation, and testing.

168 4.2 Audio and Transcription Sourcing

169 We generalize our pipeline so that our dataset can grow to include new languages and additional
 170 data in our current languages. Currently, we source all of our sentence audio and transcription data
 171 from the Mozilla Common Voice [2] project. But nothing precludes us from using our pipeline to
 172 also support data ingestion from a variety of other large speech datasets. Examples of such datasets
 173 include The People’s Speech [12] and Facebook’s Multilingual LibriSpeech [24]. We plan to include
 174 these sources as part of our regular updates to the MSWC dataset (Sec. 7).

175 Common Voice is a large and growing crowd-sourced effort to collect speech data in both widely
 176 spoken and low-resource languages. For each language, Common Voice collects public domain
 177 text from online sources such as Wikimedia along with user-contributed sentences, and volunteers
 178 record themselves reading these sentences through a web-based API. Common Voice provides a
 179 validated subset of their crowdsourced data where at least two users have listened to each submitted
 180 sentence and affirmed that the spoken audio matches the text transcription. Our initial release of the
 181 MSWC dataset utilizes the validated subset of Common Voice version 3. Common Voice encodes
 182 each user-recorded sentence as a 48KHz MP3. For each language, we feed all sentence audio and
 183 text transcriptions into the next stage of our pipeline.

184 4.3 Forced Alignment

185 We use Montreal Forced Aligner [20] to generate per-word timing estimates from each pair of
 186 audio files and transcriptions. Forced alignment is a well-established family of techniques in speech
 187 processing for estimating the occurrence of speech events along multiple boundaries (e.g., syllables,
 188 words, or sentences). We train forced alignment from a flat start only on the Common Voice data
 189 itself, i.e., we do not rely on any external acoustic models. We use graphemes for our lexicons for
 190 each language, and alignment is performed via Baum-Welch for expectation maximization. Since

191 alignment is a compute-intensive process, requiring multiple days of CPU time for larger languages
 192 in Common Voice (e.g., English, German), we include all the generated Praat textgrid alignments as
 193 part of our dataset. Each alignment provides per-word timing estimates covering the entire sentence
 194 (i.e., not restricted to the keywords we extract), enabling speech researchers and commercial users to
 195 immediately leverage our alignments for all 50 languages. Since we cannot manually validate the
 196 timing estimates produced by forced alignment, we provide an algorithm to detect outliers in Sec. 4.6.

197 4.4 Word Extraction

198 Once the word alignments are generated for each audio file, we apply several heuristic thresholds to
 199 extract a subset of the words present as individual keyword samples. For 49 of the 50 languages in
 200 MSWC, we choose a minimum character length of 3 for performing extractions, excepting Chinese
 201 (zh-CN) which uses 2 due to the preponderance of shorter words, as discussed in Sec. 3.3. Since
 202 Chinese is non-space-delimited, we estimate word boundaries in text transcriptions via Stanford
 203 CoreNLP [18]. We also only extract words which have a minimum of five occurrences in Common
 204 Voice. Additionally, we perform text normalization and cleaning where possible (e.g., we filter
 205 quotation marks out of transcripts).

206 Each extraction is encoded in the opus file format using a 48KHz sample rate (the same sample rate
 207 as Common Voice), and is stored in a two-level nested subdirectory, where the grandparent directory
 208 is the language’s ISO code (following Common Voice’s convention) and the parent directory is the
 209 keyword. The name of each extracted audio file is the same as that of the source Common Voice audio
 210 file, enabling users to easily refer back to the source audio file, transcript, and available demographic
 211 data using Common Voice’s `validated.tsv` metadata file. If the same keyword appears in a
 212 Common Voice audio clip multiple times, each extraction is appended with a double-underscore and
 213 increasing index number (e.g., `common_voice_de_18122909__2.opus`).

214 4.5 Data Splits

215 For each language, we provide index files defining `train`, `dev`, and `test` data splits. We anticipate
 216 users will commonly operate on a small subset of keywords, and thus split per keyword rather than
 217 per language set. Each keyword retains at least one clip in each of the `train`, `dev`, and `test` splits,
 218 conforming to a target ratio of 80:10:10 as the clip count increases. Maximizing speaker diversity is
 219 prioritized when performing data splits. For each keyword, we first match and group clips by their
 220 original Common Voice source’s `client_id`. We then ensure that no known speakers appear in
 221 multiple data splits for the same keyword. Once a clip is assigned to a split, the clip’s keyword and
 222 `client_id` pairing is maintained across future releases. Gender balance is also optimized across
 223 splits. For keywords with sufficient clip counts, gender balance is prioritized for `test` and `dev` splits.

224 4.6 Estimating the Quality of Extractions via Self-Supervised Anomaly Detection

225 We empirically identify several (non-exhaustive) sources of error for
 226 automated keyword extractions: (1) mismatches in `<transcript, audio>`
 227 pairs from Common Voice, (2) incorrect word boundary estimates
 228 from forced alignment, (3) mispronounced words, or (4) loud back-
 229 ground noise. While existing speech recognition tools may be capable
 230 of detecting outlier samples in some languages, we recognize the need
 231 to provide sample quality estimates across all languages in our dataset.

232 We estimate whether a keyword sample is an outlier via clustering and
 233 nearest-neighbors, similar to prior work in self-supervised anomaly
 234 detection for computer vision [27, 29, 28]. For a given keyword, we
 235 randomly choose a small number of samples (e.g., 50), referred to as
 236 the *training set* in the remainder of the section. We construct a feature
 237 vector representation for each sample and cluster these features using
 238 k -means. For the remaining samples, we calculate the euclidean
 239 distance of the feature vector representation to each of the cluster
 240 centers, and choose the smallest distance as our outlier metric.

Word	# Clips	Near	Far
may	7551	4%	32%
did	14353	32%	32%
shirt	1299	0%	54%
taken	1734	0%	12%
watch	1476	4%	28%
entire	1354	0%	12%
nature	1018	0%	12%
reading	1594	8%	58%
outside	2075	0%	2%
current	1836	10%	4%
followed	1307	0%	10%
provided	1627	0%	4%
political	1474	0%	8%

Table 5: Word error rates (WER) by outlier metric.

241 For the feature representation, we use an embedding model trained on our dataset [19] which extracts
 242 a 1024-dimensional feature vector representation from a 49x40 spectrogram input. We empirically

243 select $k = 5$ as the number of clusters,² and postulate that this is sufficiently large to capture desired
 244 speaker diversity in feature space (namely, capturing variance in genders, accents, and age as inliers)
 245 while remaining small enough to reduce the likelihood of assigning a cluster to any outliers which
 246 may be present in the random training set. The larger the value of the outlier metric, the more likely
 247 we estimate the sample to be an error. This enables our dataset users to choose a threshold for
 248 discarding $n\%$ of the keyword samples with the largest outlier metric values. We provide users of our
 249 dataset with the ability to construct this outlier metric for any sample in our dataset by releasing our
 250 anomaly detection code and model.

251 To assess the efficacy of our outlier metric, we conduct listening tests on 75 randomly chosen
 252 keywords across English, Spanish, and German, and report sample results for English in Table 5 (full
 253 results in Appendix 1.1). For each keyword, we list the number of extractions in our dataset (# Clips)
 254 and among each we choose 50 at random to cluster ($k = 5$) via our feature extractor. We then sort the
 255 remaining clips by minimum distance to each cluster in feature space, and report the Word Error Rate
 256 (WER) as a percentage when listening to the nearest 50 and farthest 50 samples in the sorted dataset.
 257 Table 5 shows a clear trend that the word error rates of the nearest 50 samples for each keyword (with
 258 an outlier metric mean \pm stdev value of 1.67 ± 0.19) is significantly smaller than the word error rates
 259 of the farthest 50 samples (with a correspondingly higher outlier metric mean \pm stdev of 3.96 ± 0.3).
 260 We also evaluate 25 randomly chosen words each in Spanish and German using the above approach.
 261 For the closest 50 samples, we observe an average WER of 4.6% in Spanish and 1.04% in German,
 262 whereas for the farthest 50 samples we estimate an average WER of 21.4% in Spanish and 29.6% in
 263 German. Our detailed results for Spanish and German are provided in Appendix Table 9.

264 We note several characteristics of outliers observed in our dataset. Forced alignment is likelier to
 265 struggle to accurately capture word boundaries for shorter words (Table 5 and Appendix Table 9 are
 266 sorted by character length) and outliers often consist of a portion of a one-syllable word, which leads
 267 to higher WER across all distances on short (3 and 4 character) words. Conversely, alignment-based
 268 extractions for long, multisyllabic words (e.g., 9 and 10 character words) often remain discernible
 269 across all distances, resulting in a lower overall WER where only true outliers are filtered (e.g.,
 270 missing words in the original recording). Our distance metric is most informative for medium-length
 271 words, where errors consist both of true anomalies and misplaced word boundaries from forced
 272 alignment, for example with the words *shirt* and *watch*. Lastly, with 50 training samples and 5
 273 clusters, training outliers can occasionally be assigned clusters (e.g., with *did* and *current*), and in
 274 practice, users of our dataset may wish to tune these hyperparameters.

275 5 Evaluating Multilingual Spoken Words Corpus

276 We evaluate the ability to train keyword spotting models using our dataset. We select two model ar-
 277 chitectures and assess the top-1 accuracy of these models trained on MSWC extractions. Furthermore,
 278 we compare these accuracies to the same models trained on the current standard for keyword spotting
 279 research, Google’s Speech Commands dataset (GSC) [32]. In particular, we seek to determine if a
 280 domain gap exists between manually recorded keywords (GSC) and extracted keywords (MSWC),
 281 spoken in the context of a full sentence, where coarticulation effects may alter pronunciation.

282 To investigate this domain gap, we do the following: (1) We cross-compare the test accuracy of two
 283 DSCNN models [4], one trained on five keywords chosen from the GSC dataset and one trained
 284 on the same five keywords from MSWC and then assess the GSC-trained model’s classification
 285 performance on MSWC data and vice versa, (2) repeat the first experiment but filter the MSWC data
 286 via Sec. 4.6’s outlier metric and (3) fine-tune single-target models following [19] using only five
 287 randomly chosen samples per keyword.

288 We select five target keywords from GSC based on the number of
 289 samples available for comparison in MSWC. Table 6 shows the total
 290 number of samples available in both GSC and MSC for these five
 291 keywords. Our dataset reflects word frequencies in natural speech, and
 292 consequently exhibits imbalanced classes compared to GSC, a manually
 293 collected dataset. We therefore choose *left*, *right*, *down*, *yes*, and *off* as
 294 the number of samples for these 5 words in GSC approximately matches

Table 6: GSC targets.

keyword	# GSC	# MSWC
left	3801	5575
right	3778	7583
down	3917	8560
yes	4044	3402
off	3745	6486

²We maximized the number of anomalous samples found in the farthest 2% percentile in listening tests on 8 English words distinct from the 25 words later used to evaluate our metric’s performance.

Table 7: Domain gap between our MSWC extracted keywords and manually recorded Google Speech Commands (GSC). Rows indicate the testing dataset; columns indicate the training dataset.

(a) 5 Target DSCNN models.				(b) DSCNN on MSWC inliers.				(c) Single Target 5-Shot models.			
		Train				Train				Train	
		MSWC	GSC			MSWC	GSC			MSWC	GSC
Test	MSWC	85.3%	60.4%	Test	MSWC	88.0%	62.3%	Test	MSWC	94.7%±3.9	87.6%±7.4
	GSC	78.6%	85.2%		GSC	78.4%	85.2%		GSC	85.2%±7.4	89.0%±3.7

295 our dataset. We use an 80:10:10 ratio for train/val/test splits, following GSC’s published splits [32]
 296 and MSWC’s splits defined in Sec. 4.5. For all evaluations, we use a 49x40 input spectrogram
 297 generated by TensorFlow’s Microfrontend [31] on 1-sec 16KHz wav encodings of GSC and MSWC
 298 clips. We refer to [4] for all other DSCNN hyperparameters and [19] for hyperparameters used in
 299 5-shot transfer learning.

300 We report baseline results of our top-1 classification accuracies in Tables 7a, 7b, and 7c. In Table 7a,
 301 we use the full test split of GSC and MSWC for each keyword, combined with random samples of
 302 background audio and unknown samples each equal to the average number of samples per target
 303 keyword. Unknown samples are drawn uniformly from non-target English words in GSC and MSWC.
 304 We observe a drop in accuracy between a DSCNN model trained and evaluated on GSC data (85.2%)
 305 and a model trained on MSWC data and evaluated on the same GSC data (78.6%). The relatively
 306 small size of the measured domain gap suggests extracted keywords hold promise for utility in
 307 keyword spotting applications. Future work will seek to close this gap via domain adaptation.

308 Furthermore, we observe potential evidence in Table 7a that our dataset can provide additional
 309 robustness. We note that the relatively low accuracy (60.4%) of the 5-target model trained on GSC
 310 and tested on MSWC is due primarily to misclassifying target keywords as ‘unknown’. This behavior
 311 is likely due to the fact that the GSC model is trained only on high quality, manually recorded target
 312 samples, and therefore is more likely to classify many samples in a wider distribution as ‘unknown’.
 313 Models trained on smaller curated KWS datasets may therefore be brittle in practice.

314 Table 7b reports results for the same setup as Table 7a, after selecting the closest 80% of samples
 315 from MSWC using our outlier metric (Sec. 4.6) with $k = 5$ and 50 samples. The accuracy increase
 316 for models trained on MSWC and GSC on MSWC test data show that our outlier metric discards
 317 anomalous samples from crowdsourced data and boundary errors in forced alignment estimates.

318 Table 7c aggregates the mean and standard deviation of classification accuracies for 5-shot single-
 319 target models for each of the 5 target keywords in Table 6. Each model has three classes: *target*
 320 *keyword*, *unknown*, and *silence/background-noise*. For each target keyword, we use 5 random seeds
 321 where each random seed corresponds to a different selection of five 1-second training samples per
 322 target model (hence, Table 7c aggregates 25 models fine-tuned on MSWC target data and 25 on GSC
 323 data). Our embedding model is pretrained on our dataset to classify 760 keywords across 9 languages,
 324 and fine-tuned on the five training examples as described in [19]. Importantly, the keywords in
 325 Table 7c were not observed during pretraining of the embedding model, ensuring our few-shot results
 326 are representative for arbitrary keywords. Our results indicate that our dataset can be used to achieve
 327 high accuracy KWS models through fine-tuning, and the relatively small number of samples for
 328 keywords in our low resource languages (Table 2) is not a barrier to producing accurate KWS models.

329 6 Downstream Applications and Broader Impact

330 **Motivation, Ethics, and Biases:** We provide useful, free, and open spoken word data in under-
 331 resourced languages. Datasets of this nature are instrumental in democratizing speech technology,
 332 expanding the inclusivity of research, and widening the reach of voice applications. Here, we consider
 333 potential biases in our dataset. For our initial release, MSWC is built exclusively on pre-existing
 334 Common Voice data, thus, biases in Common Voice, such as gender or accent bias, can propagate into
 335 MSWC. Prior work has explored demographic biases in Common Voice [23], but as MSWC expands
 336 to additional datasets, we must be mindful of the biases contained in those corpora. In addition,
 337 MSWC may contain biases which stem from forced alignment. Alignments may be of lower quality
 338 across low resource languages and for shorter words [10]. Generally speaking, as the size of the

339 dataset grows, the quality of alignments generated by the Montreal Forced Aligner will increase. Our
340 (optional) outlier metric (Sec. 4.6) may inadvertently introduce additional bias - for example, too few
341 clusters may exclude some accents or non-typical speech. We report estimated errors in English, but
342 leave other languages to future analysis. We provide a Dataset Datasheet [13] in the supplementary
343 material that describes the collection, distribution, and recommended uses of MSWC.

344 **Applications:** MSWC is in use for Sustainable Development Goals in collaboration with Makerere
345 University. Public sentiment can aid in understanding the impact and reach of public health
346 interventions—in particular, radio call-in shows often reflect sentiment in low-internet connectivity
347 areas. Tools for automated radio monitoring of limited vocabulary terms help healthcare providers
348 and organizations such as the United Nations find relevant discussions to aid in decision support [22].

349 To aid with COVID-19 response, we have used MSWC to develop a deployable tool for monitoring
350 Ugandan public radio broadcasts for COVID-19 keywords in the Luganda language. As a low-
351 resource language, it is cost-prohibitive to record thousands of keyword samples for each term of
352 interest. Hence, we use MSWC to pre-train a large multilingual embedding model, and fine tune
353 keyword search models using only five to ten example recordings in Luganda. We achieve an average
354 58% true positive rate across five COVID-19 terms in Luganda in our initial deployment candidate.

355 KWS datasets are instrumental to the development of many other applications. The data can be used
356 as the target word of a model [5], or to populate the ‘other’ category of non-target keywords. When
357 targeting low-resource languages, the data can pre-train a model to achieve better accuracy with
358 fewer samples in the target language. Other use cases include wake words for virtual assistants [30],
359 keyword search [21], and voice interfaces for embedded devices [4].

360 7 Long-term Support and Future Work

361 In order for MSWC to serve as a long-term global resource to all users, we are collaborating with
362 MLCommons (mlcommons.org), a non-profit organization that supports machine learning innovation
363 to benefit everyone. We strive to meet four criteria: (1) Relevance - providing updates and expanding
364 the dataset, (2) Accuracy - fixing errors, e.g., erroneous alignments, (3) Usability - offering friendly
365 licensing terms for research and commercial applications, and (4) Accessibility - freely downloadable.

366 We will keep the dataset updated using our automated and scalable data engineering pipeline (Sec. 4)
367 and generate future versions of MSWC, given that Common Voice performs public releases regularly.
368 Per release, we will make the entire corpus available, including alignments, audio clips, and dataset
369 splits, analogous to our initial release. Each release will bring improvements to the data engineering
370 pipeline, such as improved alignment quality. We support TensorFlow Datasets (TFDS) [1], enabling
371 users to import our dataset in one line of code. We plan to expand the contents of MSWC beyond
372 Common Voice as our system can ingest other audio and transcription sources (Section 4.2).

373 8 Conclusion

374 Voice interfaces hold promise to democratize access to technology, but the lack of large multilingual
375 datasets has prevented their proliferation. We present a spoken word corpus in 50 languages with
376 over 23 million examples, as a key step in achieving global reach for speech technology. For many
377 languages, our corpus is the first available keyword dataset. We detail our dataset’s contents, and
378 report comparable accuracies on models trained using our corpus relative to prior datasets. During
379 review, MSWC is available via the private URL provided in OpenReview. Following the review period,
380 our dataset will be hosted, maintained, and advanced by the MLCommons non-profit organization.

381 Acknowledgements

382 We acknowledge the help of the individuals who contributed and the organizations that supported this
383 project. We recognize Dr. Joyce Nakatumba-Nabende and Jonathan Mukiibi at Makerere University
384 for their collaboration on our radio monitoring work. We thank Tejas Prabhune from the Evergreen
385 Valley High School in California (USA) for his work on our initial data quality analysis. Sharad
386 Chitlangia is a student of BITS Pilani (India) and did work as an intern at Harvard University. This
387 work was sponsored in part by the Semiconductor Research Corporation (SRC) and Google.

388 **References**

- 389 [1] TensorFlow Datasets, a collection of ready-to-use datasets. [https://www.tensorflow.org/](https://www.tensorflow.org/datasets)
390 datasets.
- 391 [2] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders,
392 F. M. Tyers, and G. Weber. Common voice: A massively-multilingual speech corpus. *arXiv*
393 *preprint arXiv:1912.06670*, 2019.
- 394 [3] A. Awasthi, K. Kilgour, and H. Rom. Teaching Keyword Spotters to Spot New Keywords
395 with Limited Examples. In *Proc. Interspeech 2021*, pages 4254–4258, 2021. doi: 10.21437/
396 Interspeech.2021-1395.
- 397 [4] C. Banbury, V. J. Reddi, P. Torelli, J. Holleman, N. Jeffries, C. Kiraly, P. Montino, D. Kanter,
398 S. Ahmed, D. Pau, et al. Mlperf tiny benchmark. *arXiv preprint arXiv:2106.07597*, 2021.
- 399 [5] C. Banbury, C. Zhou, I. Fedorov, R. Matas, U. Thakker, D. Gope, V. Janapa Reddi, M. Mattina,
400 and P. Whatmough. Micronets: Neural network architectures for deploying tinyml applications
401 on commodity microcontrollers. *Proceedings of Machine Learning and Systems*, 3, 2021.
- 402 [6] A. Berg, M. O’Connor, and M. T. Cruz. Keyword transformer: A self-attention model for
403 keyword spotting. *arXiv preprint arXiv:2104.00769*, 2021.
- 404 [7] G. Chen, C. Parada, and G. Heigold. Small-footprint keyword spotting using deep neural
405 networks. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing*
406 *(ICASSP)*, pages 4087–4091. IEEE, 2014.
- 407 [8] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott,
408 L. Zettlemoyer, and V. Stoyanov. Unsupervised cross-lingual representation learning at scale.
409 *arXiv preprint arXiv:1911.02116*, 2019.
- 410 [9] A. Coucke, M. Chlieh, T. Gisselbrecht, D. Leroy, M. Poumeyrol, and T. Lavril. Efficient keyword
411 spotting using dilated convolutions and gating. In *ICASSP 2019-2019 IEEE International*
412 *Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6351–6355. IEEE,
413 2019.
- 414 [10] C. DiCanio, H. Nam, D. H. Whalen, H. Timothy Bunnell, J. D. Amith, and R. C. García. Using
415 automatic alignment to analyze endangered language data: Testing the viability of untrained
416 alignment. *The Journal of the Acoustical Society of America*, 134(3):2235–2246, 2013.
- 417 [11] D. Eberhard, G. Simons, and C. Fennig. Ethnologue: Languages of the world. *SIL International*,
418 2021. URL <http://www.ethnologue.com>.
- 419 [12] D. Galvez, G. Damos, J. M. C. Torres, K. Achorn, A. Gopi, D. Kanter, M. Lam, M. Mazumder,
420 and V. J. Reddi. The people’s speech: A large-scale diverse english speech recognition dataset
421 for commercial usage. 2021.
- 422 [13] T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. Daumé III, and
423 K. Crawford. Datasheets for datasets. *arXiv preprint arXiv:1803.09010*, 2018.
- 424 [14] J. Huang, W. Gharbieh, H. S. Shim, and E. Kim. Query-by-example keyword spotting system
425 using multi-head attention and soft-triple loss. In *ICASSP 2021-2021 IEEE International*
426 *Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6858–6862. IEEE,
427 2021.
- 428 [15] B. Kim, M. Lee, J. Lee, Y. Kim, and K. Hwang. Query-by-example on-device keyword spotting,
429 2019.
- 430 [16] Kristina, D. Klein, C. Manning, and Y. Singer. Feature-rich part-of-speech tagging with a cyclic
431 dependency network. *HLT-NAACL*, pages 252–259, 2003.
- 432 [17] J. Lin, K. Kilgour, D. Roblek, and M. Sharifi. Training keyword spotters with limited and
433 synthesized speech data. In *ICASSP 2020-2020 IEEE International Conference on Acoustics,*
434 *Speech and Signal Processing (ICASSP)*, pages 7474–7478. IEEE, 2020.

- 435 [18] C. D. Manning, M. Surdeanu, J. Bauer, J. R. Finkel, S. Bethard, and D. McClosky. The stanford
436 corenlp natural language processing toolkit. In *Proceedings of 52nd annual meeting of the*
437 *association for computational linguistics: system demonstrations*, pages 55–60, 2014.
- 438 [19] M. Mazumder, C. Banbury, J. Meyer, P. Warden, and V. J. Reddi. Few-shot keyword spotting in
439 any language. *Proc. Interspeech 2021*, 2021.
- 440 [20] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger. Montreal forced aligner
441 [computer program], 2017.
- 442 [21] R. Menon, A. Saeb, H. Cameron, W. Kibira, J. Quinn, and T. Niesler. Radio-browsing for
443 developmental monitoring in uganda. In *2017 IEEE International Conference on Acoustics,*
444 *Speech and Signal Processing (ICASSP)*, pages 5795–5799. IEEE, 2017.
- 445 [22] R. Menon, H. Kamper, E. van der Westhuizen, J. Quinn, and T. Niesler. Feature Exploration
446 for Almost Zero-Resource ASR-Free Keyword Spotting Using a Multilingual Bottleneck Ex-
447 tractor and Correspondence Autoencoders. In *Proc. Interspeech 2019*, pages 3475–3479, 2019.
448 doi: 10.21437/Interspeech.2019-1665. URL [http://dx.doi.org/10.21437/Interspeech.](http://dx.doi.org/10.21437/Interspeech.2019-1665)
449 [2019-1665](http://dx.doi.org/10.21437/Interspeech.2019-1665).
- 450 [23] J. Meyer, L. Rauchenstein, J. D. Eisenberg, and N. Howell. Artie bias corpus: An open
451 dataset for detecting demographic bias in speech applications. In *Proceedings of the 12th*
452 *Language Resources and Evaluation Conference*, pages 6462–6468, Marseille, France, May
453 2020. European Language Resources Association. ISBN 979-10-95546-34-4. URL [https:](https://aclanthology.org/2020.lrec-1.796)
454 [//aclanthology.org/2020.lrec-1.796](https://aclanthology.org/2020.lrec-1.796).
- 455 [24] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, and R. Collobert. Mls: A large-scale multilingual
456 dataset for speech research. *Interspeech 2020*, Oct 2020. doi: 10.21437/interspeech.2020-2826.
457 URL <http://dx.doi.org/10.21437/Interspeech.2020-2826>.
- 458 [25] P. Qi, Y. Zhang, Y. Zhang, J. Bolton, and C. D. Manning. Stanza: A python natural language
459 processing toolkit for many human languages. *arXiv preprint arXiv:2003.07082*, 2020.
- 460 [26] C. K. Reddy, E. Beyrami, J. Pool, R. Cutler, S. Srinivasan, and J. Gehrke. A scalable noisy
461 speech dataset and online subjective test framework. *Proc. Interspeech 2019*, pages 1816–1820,
462 2019.
- 463 [27] K. Roth, L. Pemula, J. Zepeda, B. Schölkopf, T. Brox, and P. Gehler. Towards total recall in
464 industrial anomaly detection, 2021.
- 465 [28] V. Schwag, M. Chiang, and P. Mittal. Ssd: A unified framework for self-supervised outlier
466 detection, 2021.
- 467 [29] K. Sohn, C.-L. Li, J. Yoon, M. Jin, and T. Pfister. Learning and evaluating representations for
468 deep one-class classification. In *International Conference on Learning Representations*, 2021.
469 URL <https://openreview.net/forum?id=HCSgyPUfeDj>.
- 470 [30] S. Team. Hey siri: An on-device dnn-powered voice trigger for apple’s personal assistant. *Apple*
471 *Machine Learning Journal*, 1(6), 2017.
- 472 [31] TensorFlow. TensorFlow Micro Frontend. [https://github.com/tensorflow/](https://github.com/tensorflow/tensorflow/tree/master/tensorflow/lite/experimental/microfrontend)
473 [tensorflow/tree/master/tensorflow/lite/experimental/microfrontend](https://github.com/tensorflow/tensorflow/tree/master/tensorflow/lite/experimental/microfrontend), 2021.
474 [Online; accessed 25-Aug-2021].
- 475 [32] P. Warden. Speech commands: A dataset for limited-vocabulary speech recognition. *arXiv*
476 *preprint arXiv:1804.03209*, 2018.
- 477 [33] A. Werchaniak, R. B. Chicote, Y. Mishchenko, J. Droppo, J. Condal, P. Liu, and A. Shah.
478 Exploring the application of synthetic audio in training keyword spotters. In *ICASSP 2021-2021*
479 *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages
480 7993–7996. IEEE, 2021.
- 481 [34] Y. Zhang, N. Suda, L. Lai, and V. Chandra. Hello edge: Keyword spotting on microcontrollers.
482 *arXiv preprint arXiv:1711.07128*, 2017.

483 **Checklist**

- 484 1. For all authors...
- 485 (1) Do the main claims made in the abstract and introduction accurately reflect the paper's
486 contributions and scope? [Yes]
- 487 (2) Did you describe the limitations of your work? [Yes] in section 5
- 488 (3) Did you discuss any potential negative societal impacts of your work? [Yes] in section
489 6
- 490 (4) Have you read the ethics review guidelines and ensured that your paper conforms to
491 them? [Yes]
- 492 2. If you are including theoretical results...
- 493 (1) Did you state the full set of assumptions of all theoretical results? [N/A]
- 494 (2) Did you include complete proofs of all theoretical results? [N/A]
- 495 3. If you ran experiments (e.g. for benchmarks)...
- 496 (1) Did you include the code, data, and instructions needed to reproduce the main ex-
497 perimental results (either in the supplemental material or as a URL)? [Yes] the
498 code for all experiments is available here: [https://github.com/harvard-edge/
499 multilingual_kws](https://github.com/harvard-edge/multilingual_kws)
- 500 (2) Did you specify all the training details (e.g., data splits, hyperparameters, how they
501 were chosen)? [Yes]
- 502 (3) Did you report error bars (e.g., with respect to the random seed after running experi-
503 ments multiple times)? [Yes]
- 504 (4) Did you include the total amount of compute and the type of resources used (e.g., type
505 of GPUs, internal cluster, or cloud provider)? [Yes] we estimate the compute cost of
506 creating the dataset in the dataset datasheet in the supplemental materials.
- 507 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 508 (1) If your work uses existing assets, did you cite the creators? [Yes]
- 509 (2) Did you mention the license of the assets? [Yes]
- 510 (3) Did you include any new assets either in the supplemental material or as a URL? [Yes]
- 511 (4) Did you discuss whether and how consent was obtained from people whose data you're
512 using/curating? [Yes] in the dataset datasheet in the supplemental materials.
- 513 (5) Did you discuss whether the data you are using/curating contains personally identifiable
514 information or offensive content? [Yes] in the datasheet
- 515 5. If you used crowdsourcing or conducted research with human subjects...
- 516 (1) Did you include the full text of instructions given to participants and screenshots, if
517 applicable? [Yes] we include a link to the crowdsourcing website in the datasheet in
518 the supplemental materials.
- 519 (2) Did you describe any potential participant risks, with links to Institutional Review
520 Board (IRB) approvals, if applicable? [N/A]
- 521 (3) Did you include the estimated hourly wage paid to participants and the total amount
522 spent on participant compensation? [Yes] in the supplemental materials.

523 **A Appendix**

524 **1.1 Word Error Rate Analysis**

525 In Table 8, we provide the full table of word error rates from Section 4.6. To recap, we conduct
 526 hearing tests on 25 randomly chosen keywords in English to examine the efficacy of our outlier
 527 metric. We list the number of extractions in our dataset (# Clips) for each keyword choose 50 at
 528 random to cluster ($k = 5$) in feature-space using our embedding representation. The remaining clips
 529 are then expressed as embedding vectors and sorted by minimum distance to each cluster, and the
 530 Word Error Rate (WER) is reported as a percentage when listening to the nearest 50 and farthest 50
 531 samples in the sorted dataset.

532 The word error rates of the nearest 50 samples for each keyword (which have an outlier metric
 533 mean \pm stdev value of 1.67 ± 0.19) are much smaller than the word error rates of the farthest 50
 534 samples (with a correspondingly larger outlier metric mean and stdev of 3.96 ± 0.3). Using the same
 535 method, we examine 25 randomly picked words in Spanish and German and find an average WER of
 536 4.6 percent in Spanish and 1.04 percent in German for the nearest 50 samples. In comparison, we
 537 estimate an average WER of 21.4 percent in Spanish and 29.6 percent in German for the farthest 50
 538 samples. In Table 9, we show our analysis of our outlier metric’s performance on the 25 words each
 539 for German and Spanish, following the same approach described in Sec 4.6 and reiterated above.

Table 8: Word error rates (WER) by outlier metric for 25 English words.

Word	# Clips	Near WER	Far WER
may	7551	4%	32%
did	14353	32%	32%
soon	3269	6%	34%
shirt	1299	0%	54%
style	1760	0%	12%
taken	1734	0%	12%
stood	1739	0%	22%
watch	1476	4%	28%
happy	2095	2%	8%
entire	1354	0%	12%
engine	1043	0%	18%
nature	1018	0%	12%
you’ve	1886	0%	26%
reading	1594	8%	58%
village	2801	0%	34%
outside	2075	0%	2%
strange	1720	0%	18%
current	1836	10%	4%
musical	1108	0%	16%
followed	1307	0%	10%
learning	1075	2%	18%
provided	1627	0%	4%
difficult	1771	0%	8%
political	1474	0%	8%
performance	1007	6%	12%

Table 9: Word error rates (WER) by outlier metric for two additional languages. We observe an average WER of 4.6% for the closest 50 samples in Spanish and 1.04% in German, contrasted with an average WER of 21.4% for the farthest 50 samples in Spanish and 29.6% in German

(a) German				(b) Spanish			
Word	# Clips	Near WER	Far WER	Word	# Clips	Near WER	Far WER
für	25024	2%	30%	día	1110	2%	10%
was	11920	4%	32%	una	17895	32%	14%
erst	2546	4%	38%	ella	1398	8%	14%
geht	3920	0%	42%	pero	3094	4%	38%
ihre	5285	2%	26%	están	1407	8%	52%
mich	5309	4%	42%	hasta	1801	14%	24%
muss	8656	0%	44%	nació	1101	0%	14%
bitte	4093	0%	32%	algunos	1252	6%	18%
große	1221	0%	22%	carrera	1009	10%	20%
lässt	2090	0%	24%	durante	2931	0%	22%
später	3026	0%	30%	embargo	1924	0%	18%
können	6175	0%	78%	estudió	935	0%	10%
niemand	1073	0%	14%	familia	1168	0%	30%
letzten	1048	0%	44%	general	958	0%	6%
besteht	1445	0%	38%	primera	1803	2%	40%
gehören	943	0%	42%	siempre	958	0%	28%
zweiten	1244	0%	26%	entonces	1426	10%	26%
familie	1059	0%	18%	historia	1106	4%	20%
erhielt	889	0%	14%	nacional	1199	8%	26%
zunächst	1381	4%	16%	encuentra	2945	0%	22%
mitglied	858	2%	34%	principal	1024	4%	16%
gegenüber	820	4%	26%	siguiente	890	0%	26%
verwendet	1127	0%	12%	diferentes	879	0%	6%
geschichte	1097	0%	4%	universidad	1662	4%	8%
universität	1052	0%	14%	posteriormente	1120	0%	26%

540 **1.2 Extended Parts-of-Speech and Semantic Analysis in Additional Languages**

541 We extend our analysis for Parts-of-Speech tagging and semantic characterization from Section 3.4,
 542 to provide additional examples in English, Spanish, and Arabic, and to include examples for six
 543 additional languages: German, Greek, Russian, Turkish, Vietnamese, and Chinese. Our analysis code
 544 is open-sourced, enabling users to explore similar views into the other languages in our dataset. POS
 545 tags are generated using the “log-linear part-of-speech tagger” [16] in the multilingual Stanza [25]
 546 library. Zero-shot semantic characterization is performed using an NLI model based on RoBERTa [8].

Category	#K	#C	German	Category	#K	#C	Greek
Adjective	15498	515439	hauptstadt, gut, später	Noun	350	5125	βασιλόπουλο, βασιλιάς, στη
Verb	13248	666820	gibt, kommt, bitte	Verb	324	4355	είπε, είχε, ρώτησε
Noun	7805	454323	zeit, stadt, herr	Adverb	104	3196	που, στα, γιατί
Proper noun	4286	178091	zum, zur, vom	Adjective	62	723	όλα, όλοι, πρωτομάστορης
Adverb	999	359994	ein, auch, noch	Pronoun	58	4316	του, μου, την
Pronoun	162	451875	das, sie, sich	Numeral	18	592	ένα, μια, δυο
Auxiliary	96	324801	ist, sind, hat	Adposition	9	1353	στο, από, για
Other	69	1880	hei, hallo, via	Auxiliary	9	673	είναι, ήταν, έχει
Numeral	67	31045	zwei, drei, vier	Subordinating Conjunction	6	253	πώς, ότι, όταν
Adposition	44	280245	von, mit, auf	Coordinating Conjunction	5	2116	και, ούτε, αλλά
Determiner	41	468697	die, der, den	Determiner	3	488	τον, των, ενός
Punctuation	31	1555	ach, wochen, welch	Punctuation	3	183	στην, κυρ, μες
Subordinating Conjunction	14	28779	dass, wenn, weil	Particle	2	514	δεν, μην
Coordinating Conjunction	10	92660	und, aber, oder	Proper noun	1	18	γερσο
Particle	1	42287	nicht	Other	1	7	άλλα

Category	#K	#C	Russian	Category	#K	#C	Turkish
Noun	6966	203924	слово, безопасности, конференции	Noun	1567	27441	yal, avro, devam
Verb	4634	94091	является, будет, имеет	Verb	773	15161	var, değil, oldu
Adjective	3485	69481	должны, необходимо, международного	Adjective	324	8624	yeni, nasil, büyük
Adverb	437	37905	как, также, еще	Proper noun	114	2311	kosova, türkiye, srbistan
Determiner	146	21234	этой, эти, этот	Adverb	83	8505	bir, çok, daha
Numeral	86	7857	один, два, три	Pronoun	77	2772	bunun, bunu, buna
Pronoun	81	42294	что, это, его	Numeral	36	6318	iki, bin, beş
Adposition	27	8036	для, между, без	Punctuation	21	1190	yüz, kırk, nedir
Particle	16	4287	только, нет, вот	Adposition	15	1765	için, kadar, olarak
Subordinating Conjunction	8	2650	чтобы, если, когда	Coordinating Conjunction	12	1712	ancak, fakat, ile
Interjection	7	579	при, аль, ая	NaN	10	113	nelerdir, zamandır, vardır
Proper noun	5	137	firefox, microsoft, mozilla	Auxiliary	9	721	ise, musunuz, misiniz
Auxiliary	4	584	будем, буду, будущи	Determiner	7	826	her, bazı, tüm
Coordinating Conjunction	4	1104	или, либо, причем	Interjection	1	23	hey

Category	#K	#C	Vietnamese	Category	#K	#C	Chinese (zh-CN)
Noun	15	118	con, ng, i, ông	Proper noun	37	884	奥地利, 巴伐利亚州, 马来西亚
Other	7	78	không, c, còn	Noun	34	525	俱乐部, 美术馆, 篮球队
Proper noun	6	32	âu, ây, v y	Numeral	10	109	十八年, 四十四, 三尖杉
Verb	6	50	nói, tao, nhìn	Verb	6	52	影响力, 意味着, 发源地
Coordinating Conjunction	5	48	thì, r i, nh	Particle	3	27	大部份, 的, 町名, 副作用
Adposition	4	33	c a, v i, cho	Adjective	1	13	连续剧
Particle	3	16	mày, th t, thòi	Adverb	1	5	日常生活
Numeral	1	11	m t				

Figure 4: Additional Parts-Of-Speech Tags in German, Greek, Russian, Turkish, Vietnamese, and Chinese with representative samples. Number of keywords (#K) and clips (#C). Sorted by #K.

547 We expand on Table 4 and provide additional data on semantic classification in English, Spanish, and
548 Arabic in Table 10, in reference to our discussion in Sec. 3.4. In Fig. 4, we provide sample data on
549 POS tags in German, Greek, Russian, Turkish, Vietnamese, and Chinese (zh-CN language code).
550 Additionally, in Fig. 5, we provide sample data on semantic keyword characterization for these six
551 additional languages, following the same approach as in Sec. 3.4.

Table 10: An expanded view of semantic keyword characterization in English, Spanish and Arabic via zero-shot multilingual NLI with representative samples. Number of clips (#C) and keywords (#K). Sorted by #K in English.

Category	# C	# K	English	#C	#K	Spanish	#C	#K	Arabic
Event	3M	1K	School, Preakness, Commanded	153K	238	Campeonato, Episodio, Eurovisión	2K	30	اهبجد بل صبح، لاق ي
Human activity	1M	820	Shooting, Prefers, Attacking	78K	211	Expandiendo, Visitando, Carpintería	1K	28	ينغذ، يرتشملا، داغتس
Location	391K	379	Home, County, Desert	213K	324	Zona, Pueblo, Marítima	333	14	اهبجد بل صبح، لاق ي
Name	327K	303	Margot, Cooney, Alvin	30K	128	Eduardo, Peter, Francisco	1K	24	يردءا، بقعبءاس، نانا
Animal	131K	281	Sheep, Camel, Muzzle	99K	137	Águila, Especies, Jaguar	1K	25	نولءاب، ءة حافت، برف
Number	301K	279	Second, Six, Stringent	31K	75	Multiplicar, Primero, Estadísticos	493	17	ءة ثلاث، سدحاو، قرشم
General reference	113K	220	Understand, Generally, Trivial	216K	319	Recomendaciones, Recompensa, Frecuencias	185	8	ببعءلا، نوناقءلا، ءة قيرط
Common words	218K	217	Often, Popular, Usually	56K	172	Tambien, Ademas, Biográfico	206	7	ءطخءم، برطمت، قءيدء
City	106K	212	York, London, California	34K	94	Madrid, Berlín, Oxford	49	4	قءلبءلا، بوساكبء، نايلءاطبءا
Technology	52K	191	Gramophone, Television, Videotape	21K	106	Automotriz, Lego, Sensor	92	5	ربءءءلا، نءار، ابعءلا، زءافءءلا
Culture	52K	180	Popularize, Music, Style	15K	86	Tradicional, Concertista, Cantante	57	4	قوف، قءراهمب، بءبببءلا
Language	92K	132	Arabic, English, Words	8K	69	Español, Lengua, Inglesa	98	7	ببءابء، بءببببءءلا، بءءءلا
Game	48K	131	Play, Kirby's, Football	8K	46	Deportes, Partido, Equipo	-	-	-
Political	37K	124	Impeached, Campaigns, Democrats	13K	86	Intendentes, Libertario, Poder	2K	32	ببءم، بلضفءملا، بءببببء
History	52K	119	Stories, Ancient, Roman	12K	73	Conquistas, Emperador, Históricas	257	8	اھرمء، برءءءءا، بلبء

Category	#K	#C	German	Category	#K	#C	Greek
political	756	2216421	die, der, ist	human activity	60	8145	μου, της, σου
general reference	268	144785	nur, immer, nichts	common words	40	4337	και, είχε, πως
human activity	264	221804	ich, wir, man	location	39	3017	που, στο, στα
common words	258	168175	wie, wieder, habe	political	34	4701	του, δεν, από
animal	247	212687	den, einen, einer	number	18	1147	είναι, ένα, μια
location	234	104881	unter, liegt, steht	animal	17	687	βασιλόπουλο, παρά, καράβια
society	212	111054	hat, alles, anderen	name	14	820	τον, κουλός, ποιός
event	197	108376	wird, werden, kommt	general reference	11	259	τίποτα, μην, κανένας
city	186	66360	hauptstadt, stadt, gebäude	society	8	118	κανείς, άλλες, σιωπηλά
health	184	90310	haben, sein, hast	event	8	213	φώναζε, πες, πει
name	172	52039	mein, meine, heißt	health	8	134	καλά, γερο, καρδιά
technology	162	53924	können, neue, neuen	science	5	66	γνώση, ξέρει, μάθη
social science	155	40135	zwischen, kurz, familie	technology	5	58	χειράμαξα, νέο, καινούρια
science	149	34182	weiß, wissen, universität	religion	4	57	πολύδαρος, θείος, θείου
number	138	94758	eine, einem, zwei	language	3	38	λόγια, μάθει, λέξεις
history	126	31221	einmal, jahren, jahre	game	2	16	κουτσός, κονσόλα
language	123	27136	deutschen, lang, deutsche	history	2	19	βιβλίο, διηγήθηκε, παλιά
geography	106	18763	land, gebiet, insel	city	2	37	μεγάλο, μεγάλη, σχολείο
art	104	19830	art, schön, werke	social science	2	19	σκοπό, νόημα
engineer	100	27568	kann, fährt, kühlshrank	art	1	5	σκηνή
religion	100	17252	vater, kirche, glaube	mathematic	1	28	αλυσίδα, μάθημα
culture	92	15216	europäische, musik, hören	geography	1	5	σύνορα
game	69	12931	kinder, spielen, spielt	philosophy	1	9	σοφά
mathematic	39	3193	mag, mach, zylinder	software	1	6	πρόγραμμα
software	33	3134	firefox, software, anwendung				
philosophy	21	1133	philosophie, philippinen, pfui				

Category	#K	#C	Russian	Category	#K	#C	Turkish
political	290	169316	что, это, для	political	103	23785	aneak, bin, yüz
location	183	65472	нас, хотел, время	human activity	68	9716	kişi, katıldı, ediyor
number	143	39045	имеет, уже, более	common words	66	7137	her, nasıl, avro
human activity	141	37049	усилия, человека, роль	number	62	14283	bir, iki, çok
society	137	32908	наций, организации, всех	location	52	4486	yer, ülke, ülkeden
name	112	29457	его, слово, является	name	48	4089	ise, sona, ilk
event	96	21131	будет, конференции, было	society	38	2887	durum, siyasi, durumda
health	86	17181	безопасности, быть, необходимо	history	33	1759	yıl, zaman, geri
social science	78	10953	без, процесс, характер	general reference	32	2006	yaklaşık, sadece, bunlar
animal	73	14473	кроме, такие, просто	animal	30	1948	bile, hala, reddetti
technology	70	10328	ядерного, оружия, качестве	health	22	1232	iyi, ayında, musunuz
common words	63	13958	как, также, все	technology	20	1082	yeni, proje, gelecek
general reference	57	5338	связи, сотрудничество, вместе	event	18	995	olacak, geldi, festival
history	50	4595	году, года, год	geography	16	784	büyük, civarında, dünya
mathematic	46	3505	если, задача, задачи	game	12	251	maç, çocuklar, maçı
science	46	3784	доклада, деле, доклад	science	12	213	bilgi, uzmanlar, tesis
art	35	2978	международной, выразить, создать	culture	7	234	türk, etnik, kültür
language	34	2130	речь, говорить, высказать	religion	7	126	maliyeti, dini, bedeli
software	33	2719	проекта, программы, обеспечения	software	6	159	projenin, programın, program
culture	32	2056	европейский, носят, корейской	language	6	65	küçük, dil, bölünmüş
city	31	2009	экономического, экономики, столица	art	5	116	artı, sanatçı, sanatçının
engineer	25	938	опыт, опытом, подготовки	engineer	5	46	uzman, sektörde, memur
religion	21	968	шири, исламской, верим	social science	3	31	çeşitlilik, kimseyi, eğitimi
game	19	854	играет, играть, роли	city	3	36	kent, kentın, kente
geography	12	335	линии, широкие, континента	mathematic	1	8	çözmek
philosophy	12	302	принципы, мудрость, мудрости	philosophy	1	8	düşünceye

Category	#K	#C	Vietnamese	Category	#K	#C	Chinese (zh-CN)
human activity	10	294	không, thì, c a	location	18	760	奥地利, 巴伐利亚州, 马来西亚
number	3	33	con, m t, sao	number	10	128	十八年, 四十四, 二十一
event	2	13	ch t, chuy n	human activity	8	236	俱乐部, 事务所, 影响力
location	2	16	nhà, àu	name	6	87	莎草科, 虎耳草科, 戈亚斯
animal	1	7	cái	animal	5	64	十字花科, 鹅观草, 龙胆科
language	1	15	nói	city	4	63	维也纳, 大阪府, 巴拿马
religion	1	8	trinh	health	3	29	玄参科, 蛋白质, 珍珠菜
				history	3	35	巡洋舰, 记忆体, 安土桃山时代
				culture	3	26	连续剧, 哥特式, 主题曲
				political	3	24	评论家, 主任委员, 反对派
				event	2	8	获得者, 副作用
				science	2	14	天文学家, 科进士
				technology	2	31	航空母舰, 机器人
				game	1	30	篮球队
				art	1	31	美术馆
				mathematic	1	9	大学部
				philosophy	1	21	哲学家
				religion	1	13	东正教
				society	1	6	人类学

Figure 5: Additional semantic keyword characterization in German, Greek, Russian, Turkish, Vietnamese, and Chinese via zero-shot multilingual NLI with representative samples. Number of keywords (#K) and clips (#C). Sorted by #K.