

---

# Speech Wikimedia: A 77 Language Multilingual Speech Dataset

---

Rafael Mosquera Gómez<sup>\*1</sup> Julian Eusse<sup>\*1</sup> Juan Ciro<sup>1</sup> Daniel Galvez<sup>2</sup> Ryan Hileman<sup>3</sup> Kurt Bollacker<sup>4</sup>  
David Kanter<sup>5</sup>

## Abstract

The Speech Wikimedia Dataset is a publicly available compilation of audio with transcriptions extracted from Wikimedia Commons. It includes 1780 hours (195 GB) of CC-BY-SA licensed transcribed speech from a diverse set of scenarios and speakers, in 77 different languages. Each audio file has one or more transcriptions in different languages, making this dataset suitable for training speech recognition, speech translation, and machine translation models.

## 1. Introduction

Speech research has witnessed remarkable advancements in recent years, largely driven by the availability of vast amounts of data. Tasks such as automatic speech recognition (ASR), speaker recognition, and speech translation have reached robust results even in the presence of background noise, jargon, and different accents. Nevertheless, one of the fundamental challenges in speech research is the scarcity of multilingual datasets.

In this paper, we introduce the Speech Wikimedia Dataset, a supervised dataset consisting of 1780 hours of audio files, each with one or more transcripts. The last point bears repeating: much (approximately 25%) of the dataset has multiple associated transcripts, each having its own language, which is rare in datasets sourced from the internet. The dataset is specifically curated from Wikimedia Commons to address some of the key challenges in the speech recognition, machine translation, and speech translation spaces, particularly the need for diverse multilingual data and appropriate licensing for academic and commercial usage.

The paper is organized as follows. In section 2, we describe the dataset itself and training tasks that it can support; in

---

<sup>\*</sup>Equal contribution <sup>1</sup>Factored.ai <sup>2</sup>NVIDIA <sup>3</sup>Talon Voice <sup>4</sup>Long Now Foundation <sup>5</sup>MLCommons. Correspondence to: Daniel Galvez <dgalvez@nvidia.com>.

section 3, we describe related work; in section 4, we describe some limitations of the dataset today; in section 5, we conclude with some future work that this dataset can enable.

## 2. Dataset Description

To construct the Speech Wikimedia dataset, we downloaded raw video and audio from Wikimedia Commons (“<https://commons.wikimedia.org>”), which allows only content that is “free” (wik). For our purposes, this means data under CC-BY or CC-BY-SA license, or otherwise public domain. After downloading, the data was converted to 16kHz monochannel FLAC using ffmpeg. The data is uploaded to Huggingface at <https://huggingface.co/datasets/MLCommons/speech-wikimedia>

We give statistics for three possible tasks that this dataset can be used for: speech recognition, speech translation, and machine translation.

### 2.1. Licensing Information

Given that all data is public domain, CC-BY-licensed, or CC-BY-SA licensed, we are licensing the dataset as CC-BY-SA. Following the requirements imposed by the CC-BY and CC-BY-SA licenses of our sources, accreditation is provided in the linked [credits.json](#) file.

### 2.2. Audio with Subtitles in the Same Language for Speech Recognition Task

In order to determine the amount of data available for the ASR task we used only those audios and transcriptions where the language coincided. Since Wikimedia Commons’s transcripts’ filenames contain the language of the contained text, we simply extracted the languages from these filenames. For example, the file “Elephants.Dream.ogg.en.srt” is in English, as indicated by the “.en.” substring.

Given that we didn’t initially have the audio language for each file, Whisper’s (Radford et al., 2022) language detection pipeline was used. A total of 77 different languages were detected, with English, Dutch, German, Russian, and Spanish being the most common. 69.07% of the 1780 hour dataset comprises audio-transcription pairs in the same lan-

guage. We present the number of transcribed hours of each language in Table 1 in the appendix.

### 2.3. Audio with Subtitles in Different Languages for Speech Translation Task

For speech translation, we focused on audio-transcript pairs that had different languages, which corresponds to 31% of the rest of the 1780 hours. After filtering audios and transcriptions with unknown languages, we were left with a total of 628.8 hours of audio with transcripts in different languages. We present the hours of audio for the 20 most common language pairs in Table 2 in the appendix.

### 2.4. Transcript Language Pairs (Bitexts) for Machine Translation Task

While the speech translation task relies on audio from one language with a transcription from another language, machine translation focuses on pure text translation. We find paired texts by enumerating all pairs of transcripts associated with a single audio. This is interesting in particular because approximately 10.93% of the audio files have transcripts in at least three languages.

In Table 3 in the appendix, we present total hours, number of words and occurrences of different pairs of languages for the 20 most common language pairs in the dataset.

### 2.5. Topic Distribution and Audio Content

We were also interested in determining which topics were covered across the dataset. In order to analyze this, we ran a zero shot classifier (Maiya, 2020) with labels ranging different topics, and recorded the hours of audio for each of the topics. Results are depicted in Table 4 in the Appendix section. Popular topics were current events, history, and general non-fiction references.

Based on listening to several files, we also discovered that several audios are public speeches, music, and clearly pronounced single words, probably for pronunciation dictionaries like wiktionary.

## 3. Related Work

In this section we provide an overview of previous, similar datasets.

**Mozilla Common Voice** (Ardila et al., 2020) is a CC0-licensed 17,690-hour public domain corpus of single speaker read speech in 108 languages created by volunteers. In contrast, Speech Wikimedia has much more diverse audio sources.

**Multilingual Librispeech** (Pratap et al., 2020) is a CC-BY dataset of 50,500 hours of transcribed read speech in eight

languages; 6,000 of its hours are non-English. Meanwhile, our dataset contains 77 languages, and the majority of the data is also in English.

**VoxPopuli** (Wang et al., 2021) is CC0-licensed dataset containing an unsupervised set of 400,000 hours in 23 languages, and 1,500 hours of transcribed audio in 15 languages. Like our dataset, it also contains a subset suitable for a speech translation task.

**Multilingual Spoken Words Corpus** (Mazumder et al., 2021) is a CC-BY licensed 6,000-hour dataset, containing more than 340,000 keywords in 50 different languages. It is for training keyword spotting models, not speech recognition models.

**opensubtitles** (Lison & Tiedemann, 2016) is a machine translation dataset containing 1,782 language pairings extracted from movie subtitles in 62 languages. Given the data source, it is not licensed for commercial usage. In contrast, the Speech Wikimedia Dataset has 929 language pairings from 77 languages.

## 4. Limitations

The raw data is available publicly online on Hugging Face as mentioned before; however, this data is not yet processed via forced alignment of audio to transcript and bitext word alignment for transcript to transcript, and thus not able to be used immediately in training models.

We removed all video data when converting to FLAC. In future work, this data could be helpful for a multimodal task.

While collecting this dataset, we realized that there is also a collection of audio data in Wikimedia Commons without any transcripts. We have not explored this subset and have not made it available at this time, however.

Given the small size of the dataset, we are not providing a training-test split.

## 5. Conclusions

We introduce the Speech Wikimedia Dataset, a collection of audio files with transcriptions in multiple languages extracted from Wikimedia Commons. The dataset encompasses over 1,780 hours of transcribed speech in multiple languages. The CC-BY-SA license enables commercial usage. This is the first non-read multilingual speech dataset allowing for commercial usage that we are aware of other than VoxPopuli.

## References

- Wikimedia commons:licensing. <https://commons.wikimedia.org/wiki/Commons:Licensing>. Accessed: 2023-05-23.
- Ardila, R., Branson, M., Davis, K., Henretty, M., Kohler, M., Meyer, J., Morais, R., Saunders, L., Tyers, F. M., and Weber, G. Common voice: A massively-multilingual speech corpus, 2020.
- Lison, P. and Tiedemann, J. OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pp. 923–929, Portorož, Slovenia, May 2016. European Language Resources Association (ELRA). URL <https://aclanthology.org/L16-1147>.
- Maiya, A. S. ktrain: A low-code library for augmented machine learning. *CoRR*, abs/2004.10703, 2020. URL <https://arxiv.org/abs/2004.10703>.
- Mazumder, M., Chitlangia, S., Banbury, C., Kang, Y., Ciro, J. M., Achorn, K., Galvez, D., Sabini, M., Mattson, P., Kanter, D., Damos, G., Warden, P., Meyer, J., and Reddi, V. J. Multilingual spoken words corpus. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. URL <https://openreview.net/forum?id=c20jiJ5K2H>.
- Pratap, V., Xu, Q., Sriram, A., Synnaeve, G., and Collobert, R. ISCA, oct 2020. doi: 10.21437/interspeech.2020-2826. URL <https://doi.org/10.21437/2Finterspeech.2020-2826>.
- Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., and Sutskever, I. Robust speech recognition via large-scale weak supervision, 2022.
- Wang, C., Rivière, M., Lee, A., Wu, A., Talnikar, C., Haziza, D., Williamson, M., Pino, J. M., and Dupoux, E. Voxpopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation. *CoRR*, abs/2101.00390, 2021. URL <https://arxiv.org/abs/2101.00390>.

## A. Figures

Table 1. Automatic Speech Recognition Task

LANGUAGE	HOURS OF AUDIO
ENGLISH (EN)	1488.765773
DUTCH (NL)	22.167223
GERMAN (DE)	12.658670
FRENCH (FR)	7.163889
RUSSIAN (RU)	6.985941
SPANISH (ES)	6.184720
LATIN (LA)	3.066669
POLISH (PL)	3.045028
JAPANESE (JA)	2.216300
BENGALI (BN)	2.126192
SWEDISH (SV)	1.468487
CHINESE (ZH)	1.456599
ITALIAN (IT)	1.419221
PORTUGUESE (PT)	1.344584
WELSH (CY)	1.141955
BASQUE (EU)	1.008435
HINDI (HI)	0.886795
ARABIC (AR)	0.572991
UKRAINIAN (UK)	0.441770
SLOVENIAN (SL)	0.377644
KOREAN (KO)	0.367545
HEBREW (HE)	0.238240
INDONESIAN (ID)	0.207363
THAI (TH)	0.177196
CATALAN (CA)	0.161531
GREEK (EL)	0.160628
DANISH (DA)	0.150981
PERSIAN (FA)	0.132622
VIETNAMESE (VI)	0.131922
MARATHI (MR)	0.124219
PUNJABI (PA)	0.090774
MALAYALAM (ML)	0.078354
TELUGU (TE)	0.065369
KANNADA (KN)	0.033602
HUNGARIAN (HU)	0.030055
ESTONIAN (ET)	0.029325
TURKISH (TR)	0.024743
FINNISH (FI)	0.022719
CZECH (CS)	0.021120
TELUGU (TL)	0.016138
ROMANIAN (RO)	0.015280
SLOVAK (SK)	0.000766
TAMIL (TA)	0.000364

Table 2. Speech Translation Task

AUDIO LANGUAGE	TRANSCRIPT LANGUAGE	DURATION(HOURS)
ENGLISH	SPANISH	67.115705
ENGLISH	ARABIC	43.398845
ENGLISH	FRENCH	38.163062
ENGLISH	PORTUGUESE	30.952778
ENGLISH	DUTCH	24.165356
ENGLISH	GERMAN	23.678866
ENGLISH	ITALIAN	23.442334
ENGLISH	RUSSIAN	15.557022
DUTCH	ENGLISH	14.409074
ENGLISH	POLISH	12.865772
LATIN	ENGLISH	11.722308
ENGLISH	CHINESE	11.182589
HINDI	ENGLISH	10.256298
ENGLISH	TURKISH	9.471801
ENGLISH	JAPANESE	8.782186
WELSH	ENGLISH	8.761795
ENGLISH	VIETNAMESE	6.731008
RUSSIAN	ENGLISH	6.037366
DUTCH	RUSSIAN	5.438943

Table 3. Transcript Language Pairs Statistics

LANGUAGE PAIR	TOTAL HOURS	SOURCE LANGUAGE TOKEN COUNT	TARGET LANGUAGE TOKEN COUNT	BITEXTS
ENGLISH-SPANISH	135.989042	481391.0	486965.0	629
ENGLISH-FRENCH	85.782796	262040.0	255998.0	343
ENGLISH-PORTUGUESE	57.887887	200853.0	194911.0	197
ENGLISH-RUSSIAN	55.501208	149706.0	119638.0	348
GERMAN-ENGLISH	55.449766	149499.0	168156.0	394
SPANISH-PORTUGUESE	54.185978	200486.0	191356.0	166
SPANISH-FRENCH	51.878961	178583.0	182886.0	213
ENGLISH-DUTCH	49.582888	182302.0	166567.0	164
ENGLISH-ITALIAN	47.008579	131800.0	125312.0	200
FRENCH-PORTUGUESE	38.802198	147000.0	138013.0	146
ARABIC-ENGLISH	38.239120	106115.0	136589.0	182
GERMAN-SPANISH	36.046692	110171.0	127857.0	211
ARABIC-SPANISH	34.548516	109102.0	136234.0	139
ARABIC-FRENCH	34.121227	110543.0	138088.0	134
GERMAN-FRENCH	33.843628	94528.0	111353.0	204
FRENCH-ITALIAN	33.791085	117284.0	113286.0	150
SPANISH-ITALIAN	33.368969	117633.0	109450.0	162
ARABIC-PORTUGUESE	29.675284	96408.0	113835.0	98
GERMAN-ITALIAN	28.917809	85169.0	96215.0	154
FRENCH-RUSSIAN	27.784403	80372.0	63862.0	155

Table 4. Distribution of topics and their durations

TOPIC	DURATION(HOURS)
CURRENT EVENTS	641.809422
OTHER	406.496021
HISTORY	154.644719
HEALTH	151.203017
GENERAL REFERENCE	114.263664
SOCIETY	95.286515
POLITICAL	46.760335
TECHNOLOGY	46.079005
NUMBER	45.406569
BUSINESS	40.940404
SCIENCE	34.944243
CULTURE	27.520957
LANGUAGES	26.272240
CITY	24.718620
LOCATION	15.258170
SOFTWARE	8.970613
GEOGRAPHY	8.372475
ANIMAL	7.417614
RELIGION	7.382679
PHILOSOPHY	6.863642
ART	5.678405
ENTERTAINMENT	5.076250
MATHEMATICS	2.186313
CRYPTO	1.731447
GAMING	1.252531
ENGINEERING	0.154807