

Task Arithmetic with Support Languages for Low-Resource ASR

Emma Rafkin^{1,3*}

Dan DeGenaro²

Xiulin Yang¹

¹Department of Linguistics, Georgetown University

²Department of Computer Science, Georgetown University

³Applied Physics Laboratory, Johns Hopkins University
{epr41, drd92, xy236}@georgetown.edu

Abstract

The development of resource-constrained approaches to automatic speech recognition (ASR) is of great interest due to its broad applicability to many low-resource languages for which there is scant usable data. Existing approaches to many low-resource natural language processing tasks leverage additional data from higher-resource languages that are closely related to a target low-resource language. One increasingly popular approach uses *task arithmetic* to combine models trained on different tasks to create a model for a task where there is little to no training data. In this paper, we consider training on a particular language to be a task, and we generate *task vectors* by fine-tuning variants of the Whisper ASR system. For pairings of high- and low-resource languages, we merge task vectors via a linear combination, optimizing the weights of the linear combination on the downstream word error rate on the low-resource target language’s validation set. We find that this approach consistently improves performance on the target languages.¹

1 Introduction

The 2025 Mozilla Common Voice Spontaneous Speech ASR shared task focuses on developing ASR systems for 26 low-resource languages. Training data is provided for 21 languages, while the remaining five are treated as unseen. Unlike typical ASR benchmarks based on scripted speech, this task emphasizes spontaneous speech. Performance is evaluated using Word Error Rate (WER) under four tracks: overall multilingual performance (Task 1), language-specific improvement over a baseline (Task 2), baseline improvement under a 500MB model size constraint (Task 3), and generalization to unseen languages (Task 4).

*This research was completed through Georgetown and was not funded by APL.

¹Code for this project can be found at <https://github.com/ddegenaro/mozilla-asr-challenge>.

In this paper, we explore the use of task arithmetic (Nagasawa et al., 2025) to leverage higher-resource languages for low-resource ASR. We fully fine-tune the pre-trained whisper-tiny model (Radford et al., 2023) and train Low-Rank Adapters (LoRA) (Hu et al., 2022) for whisper-large-v3 for each target language. To improve performance in low-resource settings, we incorporate information from genetically related higher-resource “support languages.” For each target language, we identify closely related languages and train corresponding models on Common Voice (CV) data (Ardila et al., 2020) using the same training setup. We then apply task arithmetic to combine target and support models (or adapters), selecting merged models based on Word Error Rate (WER) on the target development sets. Our results show consistent improvements across languages, demonstrating the effectiveness of support language integration for low-resource ASR with task arithmetic.

2 Related Work

2.1 Automatic Speech Recognition Models

Recent advances in neural architectures have led to rapid progress in ASR (Cui et al., 2025; Prabhavalkar et al., 2023). Speech foundation models such as wav2vec 2.0 (Baevski et al., 2020), WavLM (Chen et al., 2022), and HuBERT (Hsu et al., 2021) achieve strong performance across benchmarks with low error rates and efficient inference. Among these, OpenAI’s Whisper family has been particularly influential. Whisper models are encoder-decoder architectures trained on 680,000 hours of weakly labeled speech (Radford et al., 2023), demonstrating that large-scale, noisy supervision can support effective multilingual ASR and speech-to-English translation.

Although Whisper is trained on nearly 100 languages, its training data is dominated by high-resource languages, leading to substantial perfor-

mance disparities cross-linguistically. Nevertheless, the public release of Whisper’s model weights enables supervised fine-tuning for new languages when labeled data is available.

2.2 Low-Resource ASR

Low-resource ASR remains challenging due to the large amounts of labeled data required by neural models. Prior work has explored multilingual transfer to mitigate data scarcity, showing that leveraging related high-resource languages can improve recognition accuracy in low-resource settings (e.g., Khare et al., 2021; Ramesh et al., 2024; Su et al., 2024; Yu et al., 2023).

A complementary line of work uses task vectors (TVs) (Ilharco et al., 2023), defined as parameter differences between fine-tuned and base models, which can be composed to transfer capabilities across tasks or domains. This “task arithmetic” can also be achieved via combining LoRA layers (Zhang et al., 2023; Chitale et al., 2023; Chronopoulou et al., 2024; Li et al., 2025). Recent studies apply task arithmetic to ASR, demonstrating that combining task vectors from related higher-resource languages with target-language vectors improves performance under limited data (Ramesh et al., 2024; Su et al., 2024; Nagasawa et al., 2025).

3 Methods

3.1 Models and Training

We use LoRA layers to fine-tune whisper-large-v3, and we apply full fine-tuning to whisper-tiny to fit within the constraints of Task 2, which calls for a model that is less than 500MB in size. Due to resource constraints, we quantize whisper-large-v3 to 4-bit. We employ early stopping based on the WER on the validation set. The hyperparameters used for training can be seen in Table 1 in Appendix A.

3.2 Leveraging Related Languages

We make use of supplemental data from various languages which are genetically related to the languages being considered in this shared task. Following Nagasawa et al. 2025, we train task vectors (or LoRA adapter-based equivalents) for each support language and then combine them with the target language models at evaluation time.

Let θ be the weights of the pretrained Whisper model and θ_S be the weights of a model fine-tuned on ASR data from the support language S . We

define the TV τ_S as:

$$\tau_S = \theta_S - \theta \quad (1)$$

Let θ_T be the weights of the model trained on ASR data in the target language. We then apply τ_S to θ_T to create the final model θ_{final} :

$$\theta_{final} = \theta_T + \lambda\tau_S \quad (2)$$

Where λ is a scaling hyperparameter.

For whisper-large-v3, we apply a similar process to the respective LoRA layers. We scale the weights of the layers in the support adapter, and add them to the respective target language adapter’s weights to produce a merged adapter. We finally apply that merged adapter to the pre-trained Whisper model.

Support languages were chosen based on genetic relatedness to the target language, shared scripts, and the availability of scripted speech data in Common Voice. We fine-tune whisper-tiny and train LoRA adapters for whisper-large-v3 separately for each target language and each support language grouping. Whisper’s language detection is disabled and the language token is fixed to a *Whisper proxy language*: a pre-training language selected to maximize token overlap with the target language while sharing the same script. This initialization aims to start adaptation from a representation space as close as possible to the target language. Mappings between target languages, support languages, proxy languages, and scripts are provided in Table 2, with language codes listed in Tables 3 and 4. The selected “high-resource” familial languages were often still relatively low-resource. In such cases, we combine multiple related languages when possible. For example, for Puno Quechua, we use 13 Quechuan languages to construct the support models.²

After training all models, we tune the task-arithmetic scaling parameter λ on the validation set for each target language. Prior work has shown that performance is sensitive to the choice of λ (Li et al., 2025; Nagasawa et al., 2025). Instead of a grid search, we use Bayesian optimization with ten evaluations over the range $[0, 1]$. Preliminary experiments indicated that this range was unnecessarily large, as all optimal values of λ were below 0.5.

²We were unable to identify suitable familial languages for Papanla Totonac and Toba Qom, and therefore do not use any support languages for these targets. For Scots, the closest familial language is English; in this case, we likewise omit support languages and rely on the fact that Whisper’s pre-training is heavily dominated by English.

3.3 Data Sources and Munging

We use standard scripted speech data from CV for all support languages as well as for the five unseen target languages (Adyghe, Basaa, Kabardian, Puno Quechua, and Ushojo). For the remaining 21 target languages, we use the task-provided the spontaneous speech data.

We remove all samples flagged as problematic by CV annotators, as well as samples with zero-length audio or empty transcripts. Audio longer than 30 seconds is truncated to match Whisper’s maximum input length. Transcripts are left unchanged after truncation, which may have introduced minor audio–text misalignment. Finally, we apply the transcript-cleaning script provided by the task to remove formatting irregularities.

We leverage the votes metadata provided in CV as a proxy for transcription quality. For the spontaneous speech datasets, transcripts typically have 0–2 votes. For the scripted speech datasets (used for unseen target languages and all support languages), CV provides separate up- and down-vote counts; we compute a single vote score as the number of upvotes minus downvotes, which is generally positive. To prioritize higher-quality data, we upsample each training sample s with v votes by including $v + 1$ copies of s in the training set. This ensures that samples with no votes are retained, while higher-confidence samples are seen more frequently during training.³

4 Results

Results for full fine-tuning whisper-tiny on the development split can be seen in Figure 1. The development data for the unseen languages comes from the CV scripted speech corpus. We see improvements from incorporating the support language across all of the target languages, improving upon the target model by up to -0.092 with an average WER improvement of -0.0367. WER can be understood as the percentage of incorrect words in the transcript. Thus, depending on the performance of the original model, a 3-9% decrease in WER can indicate a large improvement in an ASR model. Compared to the task baseline results on the development split, our small model improves upon 4 of the languages: Ushojo, Toba Qom, Basaa, and Kabardian. We expect the WER to be slightly

³Due to resource constraints, we cap each language’s dataset at 100,000 samples, which only affects the vote-upsampled Luganda dataset.

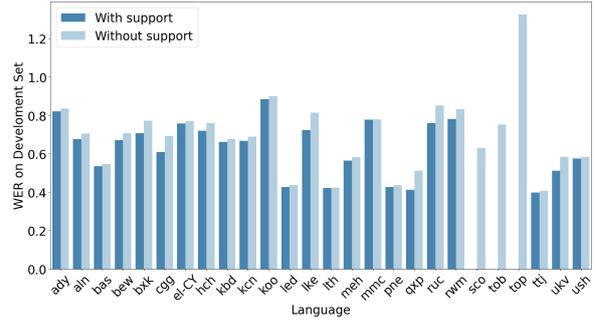


Figure 1: Incorporating the support language task vector into the whisper-tiny target language model improved all of the applicable target models.

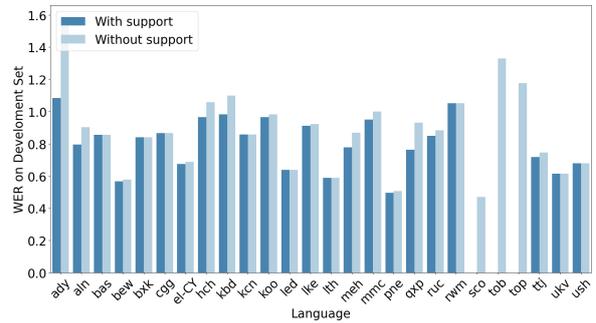


Figure 2: Incorporating the support language into the whisper-large via linear combinations of LoRA layers led to WER improvements, though overall performance is worse than whisper-tiny.

higher on the test split, specifically for the unseen language data which will be in a different domain (spontaneous speech instead of scripted speech).

Results from the whisper-large model can be found in Figure 2. For the many of languages, this model performed worse than both the task baseline and whisper-tiny, only improving upon the development baseline for Ushojo. We only submitted whisper-large transcripts for Scots, Papatla Totonac, and Betwi to task 1, as those were the only models that outperformed whisper-tiny. Additionally, there were instances in which the support language was not helpful for the target language whisper-large model. This indicates that some of these adapters were not trained well and did not appear to capture much information from the target and support languages. Despite this, the incorporation of the support language led to average WER improvements of -0.054, and in the case of Adyghe this method improved upon the base model by -0.497. This indicates that task arithmetic can be successfully performed via linear combinations of LoRA layers and is not limited to TVs.

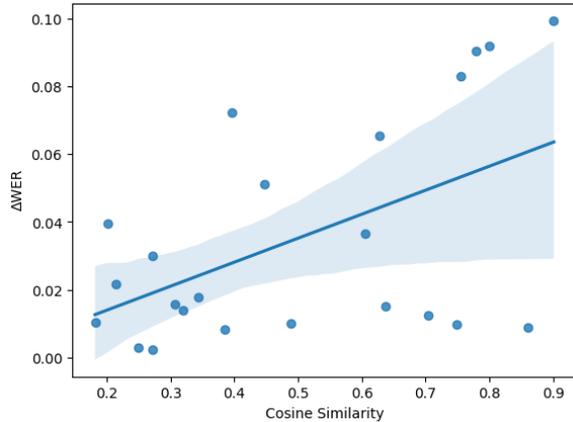


Figure 3: There is a moderate correlation between the cosine similarity of each target language and its support language and the impact of the support language model (the Δ WER) on the target model.

5 Discussion

We find that using support languages consistently improved target-language performance. In contrast to findings in the original Whisper paper (Radford et al., 2023), whisper-tiny outperformed whisper-large-v3 on many target languages. Several factors may explain this result. Due to resource constraints, we did not perform an exhaustive hyperparameter search, and the chosen LoRA rank, quantization scheme, or other training settings may have been suboptimal for the larger model. In addition, the limited number of unfrozen parameters in whisper-large-v3 may have restricted its capacity to adapt to the target languages. Or conversely, the small amount of available training data may have hindered generalization for the large model due to its higher overall capacity.

The choice of support language and the scaling parameter λ both have a substantial impact on performance. Following Li et al. (2025), we hypothesize that support languages that are more closely aligned with the target language should be more beneficial. To test this, we represent each language’s training data as a vector of token counts and compute the cosine similarity between each target language and its support language. We then examine the relationship between this similarity and the relative improvement in WER from incorporating the support language (Δ WER).

For whisper-tiny, we find a moderate positive correlation between language similarity and Δ WER, with a Pearson correlation of 0.52 ($p = 0.01$) and a Spearman correlation of 0.38 ($p =$

0.07), as shown in Figure 3. This result suggests that support languages with greater token overlap tend to yield larger improvements. Although the relatively high p value for the Spearman correlation means that this similarity metric does not reliably predict improvements in all cases, it nevertheless appears to be a reasonable heuristic for selecting support languages. We do not observe a statistically significant relationship between language similarity and the optimal value of λ .

For whisper-large-v3, we do not observe a significant correlation between language similarity and either Δ WER or λ . We suspect this is due to incomplete convergence of several LoRA adapters. Instead, we find a strong relationship between the baseline WER of the target-only model and the Δ WER, with a Spearman correlation of 0.60 ($p < 0.01$) and a Pearson correlation of 0.78 ($p < 0.01$). In other words, support languages provide greater benefit when there is more room for improvement, a pattern that is particularly evident for Adyghe. This may also help explain the smaller average gains observed for whisper-tiny, whose target-only models already perform relatively well.

Finally, the task baseline based on fine-tuning the Massively Multilingual Speech (MMS) model (Pratap et al., 2024) outperforms most of our systems. This suggests that MMS may be a more suitable base model for low-resource ASR. Additionally, it is likely that some of our choices for support languages were suboptimal, and the quantitative method for determining language similarity described above might have helped us to find better support languages.

6 Conclusion

We describe the system for our submission to the 2025 Mozilla Common Voice Spontaneous Speech challenge. We fine-tune whisper-tiny and train LoRA adapters for whisper-large on each target language and use task arithmetic with genetically related support languages sharing a script with target languages to improve upon each base model. We find that this approach consistently lowers the WER on the target languages, particularly when the target language model performs poorly alone. Rather than using the largest model possible, our results show that a more targeted approach using smaller models and relevant support data can lead to greater success for low-resource ASR.

Ethical Considerations

This approach trained 74 models in total (36 for each Whisper model type), requiring a nontrivial amount of GPU hours despite the quantization and small number of parameters being tuned per model. The environmental impact of the increasing global adoption of AI requires that researchers carefully take into consideration the compute resources needed to achieve performance, and decide what quality of model is necessary for the task at hand (Wu et al., 2022).

Additionally, many of the low-resource languages described in this paper are indigenous languages. The use and impact of research on indigenous languages is oftentimes divorced from the people that speak them. We trust that the data collection methods for these languages followed ethical codes outlined by previous research regarding indigenous language data collection (Bird, 2020; Mager et al., 2023), and hope that the impact of the systems that we explore extends beyond general “low-resource language research” and into the creation of useful technology for indigenous language speakers.

Limitations

We were unable to complete a search for the optimal training hyperparameters due to lack of a compute resources and time. In addition to training hyperparameters, the search space for λ was relatively broad: $[0, 1]$. In the end, all of the optimal values were < 0.4 . Even within the search space of $[0, 4]$, any small change in λ varied the performance significantly. For future work, this search space should be narrowed in order to allow for the optimal λ to be discovered. Furthermore, as discussed in Section 5, the choice of base model significantly impacted the overall results of this experiment. For the best performance, a better base model should be chosen and a full hyperparameter sweep should be run for language specific training values as well as λ values.

Another limitation lies in the constrained context window of our base model, Whisper, which is capped at 30 seconds. As a result, our WER scores were likely significantly negatively impacted on longer audio. Future work could involve adopting architectures without fixed-length constraints (e.g., Transducer-based models or CTC-based models like MMS (Pratap et al., 2024)) to better handle long-form recordings, which we anticipate would

further enhance the effectiveness of our approach.

References

- 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 Twelfth Language Resources and Evaluation Conference*, pages 4218–4222, Marseille, France. European Language Resources Association.
- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. [wav2vec 2.0: A framework for self-supervised learning of speech representations](#). *Advances in neural information processing systems*, 33:12449–12460.
- Steven Bird. 2020. [Decolonising speech and language technology](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3504–3519, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, and 1 others. 2022. [Wavlm: Large-scale self-supervised pre-training for full stack speech processing](#). *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1505–1518.
- Rajas Chitale, Ankit Vaidya, Aditya Kane, and Archana Santosh Ghotkar. 2023. [Task Arithmetic with LoRA for Continual Learning](#). In *Workshop on Advancing Neural Network Training: Computational Efficiency, Scalability, and Resource Optimization (WANT@ NeurIPS 2023)*.
- Alexandra Chronopoulou, Jonas Pfeiffer, Joshua Maynez, Xinyi Wang, Sebastian Ruder, and Priyanka Agrawal. 2024. [Language and Task Arithmetic with Parameter-Efficient Layers for Zero-Shot Summarization](#). In *Proceedings of the Fourth Workshop on Multilingual Representation Learning (MRL 2024)*, pages 114–126, Miami, Florida, USA. Association for Computational Linguistics.
- Wenqian Cui, Dianzhi Yu, Xiaoqi Jiao, Ziqiao Meng, Guangyan Zhang, Qichao Wang, Steven Y. Guo, and Irwin King. 2025. [Recent advances in speech language models: A survey](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13943–13970, Vienna, Austria. Association for Computational Linguistics.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. [Hubert: Self-supervised speech representation learning by masked prediction of hidden units](#). *IEEE/ACM transactions on audio, speech, and language processing*, 29:3451–3460.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-Rank Adaptation of Large Language Models](#). In *International Conference on Learning Representations*.

Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hananeh Hajishirzi, and Ali Farhadi. 2023. [Editing Models with Task Arithmetic](#). In *The Eleventh International Conference on Learning Representations*.

Shreya Khare, Ashish R Mittal, Anuj Diwan, Sunita Sarawagi, Preethi Jyothi, and Samarth Bharadwaj. 2021. [Low resource asr: The surprising effectiveness of high resource transliteration](#). In *Interspeech*, pages 1529–1533.

Hongkang Li, Yihua Zhang, Shuai Zhang, Meng Wang, Sijia Liu, and Pin-Yu Chen. 2025. [When is Task Vector Provably Effective for Model Editing? A Generalization Analysis of Nonlinear Transformers](#). In *International Conference on Learning Representations*.

Manuel Mager, Elisabeth Mager, Katharina Kann, and Ngoc Thang Vu. 2023. [Ethical considerations for machine translation of indigenous languages: Giving a voice to the speakers](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4871–4897, Toronto, Canada. Association for Computational Linguistics.

Haruki Nagasawa, Shinta Otake, and Shinji Iwata. 2025. [Task Vector Arithmetic for Low-Resource ASR](#). In *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. ISSN: 2379-190X.

Rohit Prabhavalkar, Takaaki Hori, Tara N Sainath, Ralf Schlüter, and Shinji Watanabe. 2023. [End-to-end speech recognition: A survey](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 32:325–351.

Vineel Pratap, Andros Tjandra, Bowen Shi, Paden Tomasello, Arun Babu, Sayani Kundu, Ali Elkahky, Zhaoheng Ni, Apoorv Vyas, Maryam Fazel-Zarandi, Alexei Baevski, Yossi Adi, Xiaohui Zhang, Wei-Ning Hsu, Alexis Conneau, and Michael Auli. 2024. [Scaling speech technology to 1,000+ languages](#). *J. Mach. Learn. Res.*, 25(1).

Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine Mcleavey, and Ilya Sutskever. 2023. [Robust Speech Recognition via Large-Scale Weak Supervision](#). In *Proceedings of the 40th International Conference on Machine Learning*, pages 28492–28518. PMLR. ISSN: 2640-3498.

Gowtham Ramesh, Kartik Audhkhasi, and Bhuvana Ramabhadran. 2024. [Task Vector Algebra for ASR Models](#). In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 12256–12260, Seoul, Korea, Republic of. IEEE.

Hsuan Su, Hua Farn, Fan-Yun Sun, Shang-Tse Chen, and Hung-yi Lee. 2024. [Task Arithmetic can Mitigate Synthetic-to-Real Gap in Automatic Speech Recognition](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8905–8915, Miami, Florida, USA. Association for Computational Linguistics.

Carole-Jean Wu, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, Fiona Aga, Jinshi Huang, Charles Bai, Michael Gschwind, Anurag Gupta, Myle Ott, Anastasia Melnikov, Salvatore Candido, David Brooks, Geeta Chauhan, Benjamin Lee, Hsien-Hsin Lee, and 6 others. 2022. [Sustainable ai: Environmental implications, challenges and opportunities](#). In *Proceedings of Machine Learning and Systems*, volume 4, pages 795–813.

Zhongzhi Yu, Yang Zhang, Kaizhi Qian, Cheng Wan, Yonggan Fu, Yongan Zhang, and Yingyan Celine Lin. 2023. [Master-ASR: achieving multilingual scalability and low-resource adaptation in asr with modular learning](#). In *Proceedings of the 40th International Conference on Machine Learning*, pages 40475–40487. PMLR.

Jinghan Zhang, Shiqi Chen, Junteng Liu, and Junxian He. 2023. [Composing parameter-efficient modules with arithmetic operations](#). In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23*, Red Hook, NY, USA. Curran Associates Inc.

A Hyperparameter Settings

Hyperparameter	Value
LoRA rank	32
LoRA alpha	64
LoRA dropout	0.05
LoRA target modules	q_proj, v_proj
Batch size whisper-tiny	4
Batch size whisper-large	32
Learning rate	5e-5
Maximum epochs	30
Whisper large early stopping patience	3
Whisper tiny early stopping patience	5
Gradient Accumulation	1

Table 1: Training hyperparameters

B Languages

Target Language	Language Family	Support Language	Whisper Proxy	Script
bkk	Bantu, Zone J	lg	Swahili	Latin
cgg	Bantu, Zone J	lg	Swahili	Latin
koo	Bantu, Zone J	lg	Swahili	Latin
lke	Bantu, Zone J	lg	Swahili	Latin
ruc	Bantu, Zone J	lg	Swahili	Latin
ttj	Bantu, Zone J	lg	Swahili	Latin
rwm	Bantu, Zone D	lg	Swahili	Latin
bas*	Bantu, Zone A	lg	Swahili	Latin
kcn	Afro-Asiatic, Semitic	mt	Maltese	Latin
led	Nilo-Saharan, E. Sudanic, Central	luo, kln	Somali	Latin
lth	Nilo-Saharan, E. Sudanic, W. Nilotic	luo, kln	Somali	Latin
ukv	Nilo-Saharan, E. Sudanic, E. Nilotic	luo, kln	Somali	Latin
hch	Uto-Aztecan	ncx, nhi, nlv, yaq, tar, var	Spanish	Latin
meh	Oto-Manguean, Mixtecan	cut, cux, mau	Spanish	Latin
mmc	Oto-Manguean, Otomian	cut, cux, mau	Spanish	Latin
top	Totonacan	∅	Spanish	Latin
qxp*	Quechuan	qu{p,x,r,y,s}, qx{a,u,w,t}, qv{l,a}, qw{a,s}	Spanish	Latin
tob	Guaicuruan	∅	Spanish	Latin
aln	Indo-European, Albanian	sq	Albanian	Latin
el-CY	Indo-European, Greek	el	Greek	Greek
sco	Indo-European, Germanic, West	∅	English	Latin
ush*	Indo-European, Indo-Aryan, Dardic	ur	Urdu	Perso-Arabic
ady*	Northwest Caucasian	ab	Kazakh	Cyrillic & Latin
kbd*	Northwest Caucasian	ab	Kazakh	Cyrillic & Latin
bew	Austronesian, Malay	id, ms	Malay	Latin
pne	Austronesian, Malayo-Polynesian, Kenyah	id, ms	Malay	Latin

Table 2: Low-resource target languages, their selected support counterparts and Whisper proxy language codes. * indicates a “test-only” language.

Code	Language
aln	Gheg Albanian
bew	Betawi
bvk	Bukusu
cgg	Chiga
el-CY	Cypriot Greek
hch	Wixárika
kcn	Nubi
koo	Konzo
led	Lendu
lke	Kenya
lth	Thur
meh	SW Tlaxiaco Mixtec
mmc	Michoacán Mazahua
pne	Western Penan
ruc	Ruuli
rwm	Amba
sco	Scots
tob	Toba Qom
top	Papantla Totonac
ttj	Rutoro
ukv	Kuku
ady	Adyghe
bas	Basaa
kbd	Kabardian
qxp	Puno Quechua
ush	Ushojo

Table 3: Target language codes

Code	Language
ab	Abkhaz
cut	Teutila Cuicatec
cux	Tepeuxila Cuicatec
el	Greek
id	Indonesian
klj	Kalenjin
lg	Luganda
luo	Dholuo
mau	Huautla Mazatec
ms	Malay
mt	Maltese
ncx	Central Puebla Nahuatl
nhi	Tetelancingo Nahuatl
nlv	Orizaba Nahuatl
qup	Southern Pastaza Quechua
qur	Quechua Yanahuanca
qus	Quechua Santiago del Estero
qux	Quechua Yauyos
quy	Quechua Chanka
qva	Quechua Ambo-Pasco
qvl	Quechua Cajatambo
qwa	Quechua Corongo Ancash
qws	Quechua Sihuas Ancash
qxa	Quechua Chiquián
qxt	Quechua Pasco Santa Ana de Tusi
qxu	Quechua Arequipa-La Unión
qxw	Quechua Jauja Wanka
sq	Albanian
tar	Central Tarahumara
ur	Urdu
var	Huarijio
yaq	Yaqui

Table 4: Support language codes