

# Fine-tuning Whisper on Low-Resource Languages for Real-World Applications

Vincenzo Timmel<sup>1</sup>, Claudio Paonessa<sup>2</sup>, Manfred Vogel<sup>1</sup>, Daniel Perruchoud<sup>1</sup>, Reza Kakooee<sup>1</sup>

<sup>1</sup>University of Applied Sciences and Arts Northwestern Switzerland

<sup>2</sup>Noxenum.io

{vincenzo.timmel, manfred.vogel, daniel.perruchoud, reza.kakooee}@fhnw.ch, claudio.paonessa@noxenum.io

## Abstract

This paper presents a new approach to fine-tuning OpenAI’s Whisper model for low-resource languages by introducing a novel data generation method that converts sentence-level data into a long-form corpus, using Swiss German as a case study. Non-sentence-level data, which could improve the performance of long-form audio, is difficult to obtain and often restricted by copyright laws. Our method bridges this gap by transforming more accessible sentence-level data into a format that preserves the model’s ability to handle long-form audio and perform segmentation (by predicting timestamps) without requiring non-sentence-level data. Our data generation process improves performance in several real-world applications and leads to the development of a new state-of-the-art speech-to-text (STT) model for Swiss German. We compare our model with a non-fine-tuned Whisper and previous state-of-the-art Swiss German STT models, where our new model achieves higher BLEU scores. Our results also indicate that the proposed method is adaptable to other low-resource languages, supported by written guidance and code that allows the creation of fine-tuned Whisper models, which keep segmentation capabilities and allow the transcription of longer audio files using only sentence-level data with high quality.

## 1 Introduction

Swiss German refers to the dialects spoken in the German-speaking regions of Switzerland. Due to the limited number of speakers, linguistic resources are scarce, also because Swiss German exists only as a spoken language, without any formal grammar or standardized written form. As a result, Swiss German STT systems are typically formulated as speech translation tasks, where Swiss German audio is transcribed into standard German text (Plüss

et al., 2021, 2022, 2023). Generally, "Swiss German transcription" or "Swiss German ASR" refers to converting spoken Swiss German directly into written Standard German, combining transcription and translation in a single step.

The Whisper models (Radford et al., 2022) developed by OpenAI are trained on a large-scale corpus of audio recordings and corresponding transcriptions obtained by web crawling. The dataset used for the multilingual version of Whisper includes samples from almost 100 different languages. After English and Chinese, German represents the third-largest part of the dataset, with 13’344 hours. The unexpectedly high transcription quality of Whisper for Swiss German audio and video (see Table 4) and observed distinct hallucinations (discussed in Section 5) prove the presence of Swiss German audio in the original training dataset. While the model performs remarkably well for Swiss German, there is still considerable room for improvement for practical applications that require higher transcription quality, such as judicial interrogation transcripts or medical diagnosis and treatment orders.

In addition, fine-tuning solutions which fail to predict timestamps make it impossible to use Whisper for subtitling, multi-speaker conversation pattern analysis and other real-world applications<sup>1</sup>. Table 1 illustrates how Whisper Large-v2 fine-tuned on sentences loses its capabilities to predict timestamps and starts to fail when the audio segment is getting longer and more difficult to predict, even though when evaluating it on the test-split of the sentence-level dataset, the fine-tuned model shows much improvement over the original Whisper Large-v2.

In this paper, we focus on fine-tuning OpenAI’s Whisper model for low-resource languages in real-world applications, using Swiss German as a case study. We evaluate the segmentation capabilities

<sup>1</sup><https://huggingface.co/blog/fine-tune-whisper>

Table 1: Comparison of a rapidly spoken Swiss-German saying (Straub, 2024), transcribed by Whisper Large-v2. The model, fine-tuned on sentence-level data, fails to predict timestamps and performs worse than the original Large-v2.

<b>Input Audio</b>
...
Ich zeig der, wo de Bartli de Moscht holt.
...
<b>Whisper Large-v2</b>
...
[00:00:08] Ich zeige dir, wo Bartli den Most holt. [00:00:11]
...
<b>Whisper Large-v2 (fine-tuned on sentences)</b>
...
Ich zeige dir, wo es die Bartli in den Most holt.
...

of Whisper after fine-tuning and demonstrate the beneficial effect of fine-tuning on long-form audios generated from sentence-level data. Finally, we evaluate the impact of the amount of training data on model performance when fine-tuning Whisper.

In particular, we address the following key research questions:

- How can sentence-level datasets be adapted to effectively train Whisper models for longer audio sequences while maintaining segmentation and transcription quality?
- How does fine-tuning Whisper affect its segmentation capabilities, especially when moving from sentence-level to long-form data?
- How does fine-tuning Whisper with additional datasets, such as pseudo-labeled long-form audio, affect its performance in various real-world scenarios?

By exploring these research questions, this paper provides insights into improving STT systems for low-resource languages through innovative data generation and fine-tuning strategies.

## 2 Related work

Despite the challenge of scarce data resources, recent advancements in speech translation for Swiss German have been substantial, driven in part by recent collections of high-quality sentence-level datasets from Swiss parliaments minutes and crowdsourcing initiatives such as SwissDial (Dogan-Schönberger et al., 2021), the Swiss Parliaments Corpus SPC (Plüss et al., 2021), SDS-200

(Plüss et al., 2022), and STT4SG-350 (Plüss et al., 2023).

Prior to the release of Whisper, a commonly applied foundation model for building ASR systems was XLS-R (Babu et al., 2021). XLS-R is based on the wav2vec 2.0 architecture (Baevski et al., 2020) and was pre-trained on 436K hours of speech data in 128 languages. Previous research on Swiss German speech recognition has therefore often used the pre-trained XLS-R backbone (Plüss et al., 2022, 2023; Schraner et al., 2022; Paonessa et al., 2023).

There are many papers and blog posts on the fine-tuning of Whisper (Gandhi, 2022; Singh, 2023; de Bruin, 2023; OpenAI, 2023; Ma et al., 2023; Shamsian et al., 2024; Do et al., 2023; Ferraz et al., 2023; Sicard et al., 2023) and some focused on low-resource settings (Piñeiro-Martín et al., 2024; Hsu et al., 2024; Liu et al., 2024; Pillai et al., 2024). The problem of language forgetting is also discussed extensively, and it is shown that fine-tuning on a new language yields the best performance for the new language, but degrades the capabilities on existing languages (Qian et al., 2024). However, it is rare to find papers that explicitly address the problem of fine-tuning Whisper for transcription of longer audios. Many papers also do not evaluate or discuss the segmentation capabilities after fine-tuning. Finally, it is uncommon to see evaluations of fine-tuned Whisper on out-of-distribution datasets.

## 3 Approach

### 3.1 Data Generation

The Whisper model works with a fixed input length of 30 seconds. Samples shorter than 30 seconds must be padded by appending zeros (silence). However, available datasets often consist of sentence-level samples, which are usually much shorter than 30 seconds. This is also true for Swiss German speech translation corpora such as SPC, SDS-200 and STT4SG-350. Using such training data to fine-tune Whisper is challenging because it requires significant padding for each sample. This carries the potential risk of compromising the model’s ability to robustly handle long-form audio and predict timestamps, a crucial aspect of the model for many use cases. We therefore start with available sentence-level pairs of Swiss German audio and standard German transcriptions (see Table 3) and concatenate multiple sentences to synthetically generate long-form audios with corresponding segment

timestamps, as shown in Figure 1.



Figure 1: Illustration of generated long-form training data from sentence-level audios. Although timestamps are available via the length of the audio, they are not displayed here.

The data generation strategy contains the following steps:

- **Timestamp Correction:** By using Voice Activity Detection (VAD), specifically leveraging Silero Models<sup>2</sup>, we correct the start and end timestamps of the resulting audio segments.
- **Noise Overlapping:** By simply concatenating two audio samples, the transitions often become noticeable because they abruptly change noise characteristics. To improve the transitions between consecutive samples, we employ a random overlapping technique that leverages the silence intervals detected by VAD at the beginning and end of each sample. By taking advantage of these silence parts, this enhancement accurately simulates consecutive audio segments. Together with Timestamp Correction, it also allows to create a speech overlap, such that two speakers speak over each other.
- **Speaker Retention:** For samples that include speaker identification, the probability of retaining the same speaker in successive samples is 50%. This enhancement helps to create more realistic sequences in which speaker changes occur at a pace with natural speech patterns.

Later in the study, in Table 5, we show the influence of these data generation strategies on different datasets. In Figure 2 the general approach is shown. The beginning and the end of speech are detected and then, when concatenating sentences together, we can either:

<sup>2</sup><https://github.com/snakers4/silero-models>

- do **Concat**, concatenate files as they are,
- detect the end and start of speech and make the non-speech **Overlap** by up to 200ms when concatenating audios,
- introduce a **Negative Overlap**, so that the speech of two sentences overlaps by 200ms.



Figure 2: Illustration of the logical structure for stitching together sentences using VAD and overlap mechanisms. With the help of a VAD model, we precisely mark the start and end of speech. This allows us to vary the length of pauses between sentence and even introduce an overlap.

### 3.2 Training Details

For model initialization, we use the Whisper Large-v2 weights, as initial tests showed it outperformed Whisper Large-v3 on the Swiss German datasets used. And we take advantage of its strong baseline performance for Swiss German by using the German language tag **DE** (see Section 4.2).

Using gradient check-pointing (Sohoni et al., 2022), gradient accumulation, and an 8-bit optimizer (Dettmers et al., 2021), we achieve an effective batch size of 256 on a single NVIDIA A100 40 GB GPU. Gradient check-pointing is applied to both the encoder and decoder, with 16 gradient accumulation steps and a per-step batch size of 16. Additionally, we apply stochastic depth (Huang et al., 2016) to the encoder and decoder blocks. This setup results in a mixed-precision training run that takes about 42 hours.

We use a learning rate scheduler with a linearly increasing warm-up phase followed by a linear decay to zero, as described in the original Whisper training procedure (Radford et al., 2022). During training, each sample has a 50% chance of containing timestamps and a 50% chance of containing prompts, mentioned in a comment on the OpenAI Whisper repository<sup>3</sup>.

<sup>3</sup><https://github.com/openai/Whisper/discussions/838>

Table 2: Training Hyperparameters

Parameter	Value
Optimizer	AdamW
Max. Learning Rate	$2.0 \times 10^{-4}$
Weight Decay	0.1
Warmup Updates	128
<b>AdamW Specific Parameters</b>	
$\beta_1$	0.9
$\beta_2$	0.98
$\epsilon$	$1.0 \times 10^{-9}$

Following the improved training procedure of Whisper Large-v2, we apply SpecAugment (Park et al., 2019) during training with the same parameters as in (Radford et al., 2022), summarized in Table 2.

### 3.3 Train, Validation and Test Data

For our training data, we use the Swiss German sentence-level datasets (Plüss et al., 2021, 2022, 2023) with the predefined train and validation sets. For the train and validation split, unless mentioned otherwise, we use our data-generation pipeline explained in section 3.1. As additional training and validation data, we use Swiss Broadcasting Corporation (SRG) shows, pseudo-labeled (PL) by transcription with Whisper Large-v2. We selected 17 TV series, in which Swiss German is spoken.

As test data, we use the predefined split of the Swiss German sentence-level datasets (not processed by our pipeline and thus stay as single sentences) and a Dataset-A containing a manually transcribed doctor-patient conversation obtained from a confidential phone call. Due to data privacy, this dataset cannot be disclosed and remains a closed source dataset. As an additional test set, we use SRG data from 5 TV series for which manual transcriptions are available, i.e.: *Einstein*, *Puls*, *Impact Investigativ*, *SRF Kids News*, and *SRF ohne Limit*. In contrast to the pseudo-labeled SRG train and validation data, we use as test set subtitles manually created by SWISS TXT, a subsidiary of SRG.

Because we have reasonable suspicion (see Section 5) that OpenAI has data from the SRG in its Whisper training corpus, we only considered SRG data for the validation and test set broadcasted after the release of Whisper Large-v2 to allow a fair comparison with our baseline, the Whisper Large-v2 base model.

The total hours of data used for training, validation and testing are presented in Table 3.

Table 3: Overview of the datasets used for training, validation, and testing, including totals per split.

Name (Variant)	Split	Hours	# Speakers
SDS-200 (Clean)	Train	50	1,799
STT4SG-350 (All)	Train	276	219
SPC (0.9 IOU)	Train	176	194
SRG (PL)	Train	406	–
Total		908	> 2,212
SDS-200 (Clean)	Val	5.2	288
STT4SG-350 (All)	Val	21	219
SRG (PL)	Val	20	–
Total		46.2	> 507
SDS-200 (Clean)	Test	5.2	281
STT4SG-350 (All)	Test	34	56
SPC (0.9 IOU)	Test	6	26
SRG (SWISS TXT)	Test	20	–
Dataset-A	Test	0.22	2
Total		65.42	> 365

In Figure 3, we analyze the relationship between the BLEU score (Papineni et al., 2002) on the STT4SG-350 test set and the amount of training data used for fine-tuning. For training, we used the 502 hours long-form corpus consisting of SDS-200, STT4SG-350 and SPC, but we do not include the pseudo-labeled data to show what can be expected from high-quality labeled data. The models were trained using hierarchically nested datasets, each partition holding approximately 20% of the training data. Training was run until the word error rate (WER) on the validation set showed no more improvement. Once the training had stabilized, the best performing model on WER was selected.






Figure 3: BLEU score on the STT4SG-350 test set vs. amount of training data (given in Table 3) used for fine-tuning. The model evaluated at 0 hours of training data corresponds to the original Whisper Large-v2. The SOTA model is discussed in section 4.4.

Based on the unabated rise of the curve in Figure 3, it’s reasonable to deduct that more training data will improve model performance further.

## 4 Experiments

### 4.1 Evaluation

In all experiments, the computed metrics are derived from the WhisperX<sup>4</sup> versions of the models using fp16 precision. The beam size is configured to 5, **VAD** is enabled and the language tag is **DE**. For the sentence-level datasets, we report the WER or the BLEU metric or both, if the layout allows it. For the long-form test dataset we replace the BLEU metric with SubER (Wilken et al., 2022); this allows us to incorporate a metric that also assesses the quality of segmentation by timestamp prediction. Before calculating the metrics, all sentences are transformed to lowercase, and punctuation is removed. For the BLEU metric we specifically use sacreBLEU (Post, 2018) with default parameters.

### 4.2 Base Results

In Table 4 we compare the performance of the original Whisper Large models (v2 and v3) on our datasets. Interestingly, Whisper Large-v3 exhibits an improvement over its predecessor, Large-v2, only on the STT4SG-350 test set. It is noteworthy that the 24,605 samples in the STT4SG-350 test set yielded identical transcripts for both models in 11,340 instances (46.1%). Conversely, both the SPC and SDS-200 test sets reveal a slight decline in performance for Whisper Large-v3. The gap between the two models is most evident on

the SRG dataset, where the Large-v2 model yields noticeably better performance.

Table 4: Performance of the original Whisper models on various sentence-level test sets. SubER is shown on the long-form SRG data.

Test Dataset	Model	WER	BLEU	SubER
SPC	Large-v2	<b>28.21</b>	<b>58.08</b>	-
	Large-v3	28.94	57.90	-
SDS-200	Large-v2	<b>27.69</b>	<b>57.35</b>	-
	Large-v3	27.88	57.00	-
STT4SG-350	Large-v2	22.41	<b>64.36</b>	-
	Large-v3	<b>22.01</b>	64.13	-
SRG (SWISS TXT)	Large-v2	<b>28.42</b>	<b>63.61</b>	<b>30.63</b>
	Large-v3	38.69	56.31	42.58

Based on this comparison, there is no indication of a noticeable improvement on Swiss German audio when using Whisper Large-v3 instead of Large-v2.

### 4.3 Segmentation Forgetting & Out-of-Distribution Performance Degradation

To assess the impact of fine-tuning without timestamps on both segmentation capabilities and performance on out-of-distribution data, we conducted experiments comparing models trained on different data generation parameters with our pipeline. Table 5 presents the BLEU and SubER metrics for each model across various test datasets. The following parameters for the data generation are compared:

- **Sentence-level:** Training is conducted at the sentence level with padding.
- **Concat:** No adjustments are made; sentences are simply concatenated together.
- **Speaker\_ret:** There is a 50% chance that two consecutive sentences originate from the same speaker without reusing any sentences.
- **Overlap:** There is a 50% chance that two sentences overlap, meaning the speech of the sentence begins immediately as the speech of the preceding sentence ends.
- **Neg\_overlap:** There is a 10% chance that when audios overlap, the speech of two separate sentences overlaps by 200ms, simulating scenarios where two speakers talk simultaneously.

<sup>4</sup><https://github.com/m-bain/whisperX>

- **All**: This method combines all the aforementioned preprocessing techniques.

Table 5: Comparison of BLEU and SubER metrics for Whisper Large-v2, fine-tuned on sentence-level (first row) and differently generated long-form data (Large-v2 refers to the original model without fine-tuning). The results on SRG and Dataset-A show that the sentence-level fine-tuning performs noticeably worse compared to training on long-form data.

Test Dataset	Data/Model	BLEU	SubER
STT4SG-350	Sentence-level	<b>77.38</b>	-
	Concat	76.08	-
	Speaker_ret	76.91	-
	Overlap	72.12	-
	Neg_overlap	76.53	-
	All	76.55	-
	Large-v2	64.36	-
SRG (SWISS TXT)	Sentence-level	47.57	51.34
	Concat	51.34	42.87
	Overlap	49.68	44.15
	Speaker_ret	52.51	41.07
	Neg_overlap	50.63	43.56
	All	51.62	41.64
	Large-v2	<b>63.61</b>	<b>30.63</b>
Dataset-A	Sentence-level	35.01	55.44
	Concat	46.80	41.94
	Speaker_ret	46.79	41.99
	Overlap	45.84	43.62
	Neg_overlap	45.40	43.03
	All	47.22	41.15
	Large-v2	<b>48.89</b>	<b>39.12</b>

The model fine-tuned solely on padded sentence-level samples exhibits a substantial decline in timestamp prediction accuracy, evidenced by a SubER score exceeding 51 on the SRG dataset, even with the help of VAD, which substitutes timestamps. Because we use WhisperX to evaluate the models, the timestamps are given by the VAD-Model. Without WhisperX, the metrics of the sentence-level Model would be much worse.

While the sentence-level model attains with 77.38 the highest BLEU score on the STT4SG-350 dataset, it performs poorly on longer audio sequences in SRG and Dataset-A, especially when the segmentation quality is taken into account. This underscores its limitations with out-of-distribution data. In comparison, models trained with the generated long-form dataset — especially 'All' — demonstrate better generalization, maintaining higher BLEU scores and lower SubER scores across different datasets, nearly reaching the original Whisper Large-v2 on the Dataset-A on BLEU.

The original Large-v2 model without fine-tuning, outperforms the sentence-level model on long-form

datasets. This suggests that fine-tuning exclusively on sentence-level, which was data processed into long-form, degrades the performance on unseen datasets. Incorporating long-form audio and datasets with diverse distribution into the training process is essential for preserving segmentation capabilities and ensuring robust performance across diverse data distributions.

#### 4.4 Overcome Shortcomings

Despite the measures taken to simulate long-form audio, our fine-tuning procedure leads to a reduction in segmentation and transcription quality when applied to real long-form audio, as shown by the SubER metric given in Table 5 for the Dataset-A and SRG datasets. To address this issue, we enrich the training dataset by incorporating samples from the specific distribution of the intended prediction target, in our case, pseudo-labeled SRG data mentioned in Section 3.

As part of our methodology to mitigate language forgetting in the final model training, we use the Mozilla Common Voice 13 German dataset (Ardila et al., 2020). Using the data generation pipeline described in Section 3.1, we curated a subset of 15,000 samples, each lasting 30 seconds, resulting in 125 hours of additional training data. The train, validation, and test set splits were taken as defined by Mozilla Common Voice Version 13.

This leads to a model fine-tuned on long-form audio with our generated corpus based on the three sentence-level datasets (SPC, SDS-200, and STT4SG-350), the pseudo-labeled dataset SRG (PL) based on Swiss German TV shows, and the German part of the Common Voice 13 training data, concatenated as described above. This strategy significantly improves the model performance, as shown in Table 6, and leads to a new state-of-the-art model for Swiss German speech-to-text, referred to as **SOTA**.

#### 4.5 Dialect Comparison

Since the STT4SG-350 test dataset contains identical sentences in 7 different dialects, it allows a fair comparison of model performance in terms of dialect-specific accuracy.

The results in Table 7 show large differences in the performance of the original Large-v2 model across the different dialect regions. In contrast, our fine-tuned SOTA model exhibits improved WER over a much narrower range across all dialects. For the reader, the improvements over other Swiss-

Table 6: Our new Whisper Large-v2 based SOTA model, fine-tuned on long-form audio created with data generation method 'All' and supplemented with SRG (PL) and Common Voice 13 de, compared to the original Whisper Large-v2

Test Dataset	Model	WER	SubER	BLEU
SPC	Our SOTA	<b>20.98</b>	-	<b>68.34</b>
	Large-v2	28.21	-	58.08
SDS-200	Our SOTA	<b>16.70</b>	-	<b>72.69</b>
	Large-v2	27.92	-	57.00
STT4SG-350	Our SOTA	<b>12.11</b>	-	<b>78.08</b>
	Large-v2	22.41	-	64.13
SRG (SW-TXT)	Our SOTA	<b>26.31</b>	<b>29.76</b>	<b>64.67</b>
	Large-v2	28.42	30.63	63.61
Dataset-A	Our SOTA	<b>34.50</b>	<b>35.31</b>	<b>51.40</b>
	Large-v2	38.00	39.12	48.89
CV13 de	Our SOTA	<b>6.42</b>	-	-
	Large-v2	6.53	-	-

German ASR models based on Wav2Vec (XLS-R) and Transformer (TF) architectures (Schraner et al., 2022) are also shown.

Table 7: WER for Swiss German dialects on the STT4SG-350 test set for selected models; XLS-R and TF are older models based on Wav2Vec (XLS-R) and Transformer (TF) architectures (Schraner et al., 2022).

Dialect	Large-v2	SOTA	XLS-R	TF
BS	25.02	<b>12.72</b>	16.30	21.24
BE	25.92	<b>13.68</b>	15.74	20.96
GR	19.59	<b>11.45</b>	14.32	17.29
IS	17.63	<b>10.73</b>	13.26	16.37
OS	21.27	<b>12.45</b>	16.45	18.58
VS	29.31	<b>12.72</b>	17.75	22.64
ZH	18.29	<b>11.03</b>	13.41	17.30

## 5 Conclusions

A key advantage of OpenAI’s Whisper model is its ability to process audio of arbitrary length with built-in segmentation capabilities. However, fine-tuning such a model on sentence-level datasets while preserving these features is a significant challenge.

This paper demonstrates the potential of fine-tuning Whisper for low-resource languages, using Swiss German as a case study, and addresses the three research questions posed in the introduction. First, the paper shows how sentence-level datasets can be effectively adapted for training on longer audio sequences through a novel data generation

pipeline, including techniques such as timestamp correction, noise overlapping, and speaker retention. These methods enable the generation of realistic long-form audio data that preserves segmentation and transcription quality.

Second, the fine-tuning approach significantly improves the model’s segmentation capabilities, particularly for long-form data, compared to sentence-level models. By evaluating the segmentation performance with SubER metrics, the study highlights the benefits of incorporating diverse training data and demonstrates improved robustness for timestamp prediction and audio segmentation.

Finally, the inclusion of additional datasets, such as pseudo-labeled long-form audio from Swiss Broadcasting Corporation shows, improves the model’s performance in real-world applications. We also show how to maintain performance in other languages by supplementing the training data with samples from those languages, thereby mitigating catastrophic forgetting. The results show that this method generalizes well to out-of-distribution datasets, achieving state-of-the-art performance in Swiss German speech-to-text tasks and suggesting broader applicability to other low-resource languages.

In addition, we have highlighted that while the model may improve on data from the same distribution as the training data, in reality the model performs worse on out-of-distribution data, as shown in table 5. This underscores the importance of creating or acquiring evaluation datasets that closely mimic the intended deployment environment, ensuring the ASR system’s robustness and usefulness in real-world applications.

Our research lays the groundwork for future work on data preparation and fine-tuning for OpenAI’s Whisper model, especially in low-resource settings. For this we provide a simple framework, addressing catastrophic forgetting through long-form data generation and pseudo-labeling, enabling robust transcription even with limited datasets. The code for our data generation procedure<sup>5</sup> and model fine-tuning<sup>6</sup> is publicly available.

Additionally, we observed distinct hallucinations of Whisper Large-v2 mentioning Swiss subtitling companies, such as being able to reliably trigger Whisper to transcribe "Untertitel von SWISS TXT"

<sup>5</sup><https://github.com/i4ds/Whisper-prep>

<sup>6</sup><https://github.com/i4ds/Whisper-finetune>

- a watermark of SWISS TXT that is only present in the subtitle files, never in the audio - when asked to transcribe the title music of the SRF Meteo show or when music is being played in "SRF bi de Lüt".

## 6 Future work

As we have extensively analyzed and evaluated different methods to generate long-form data from sentence-level data, the combination of the data generation and methods to avoid catastrophic forgetting, as presented by (Qian et al., 2024) by using Elastic Weight Consolidation (Kirkpatrick et al., 2017), could be a next research topic. Another potential next step involves diversifying the data sources by augmenting the pseudo-labeled datasets with additional real-world data, including a broader range of TV programs, varied conversational contexts, and noisy environments. This expansion aims to enhance the robustness and generalization capabilities of the models. Notably, preliminary experiments indicate that a fine-tuned Whisper Large-v3 model performs particularly well on conversational speech, highlighting its potential superiority in this context and emphasizing the need for a large corpus of freely spoken dialogues.

## References

- R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber. 2020. Common voice: A massively-multilingual speech corpus. In *Proceedings of the 12th Conference on Language Resources and Evaluation (LREC 2020)*, pages 4211–4215.
- Arun Babu, Changhan Wang, Andros Tjandra, Kushal Lakhotia, Qiantong Xu, Naman Goyal, Kritika Singh, Patrick von Platen, Yatharth Saraf, Juan Pino, Alexei Baevski, Alexis Conneau, and Michael Auli. 2021. XLS-R: self-supervised cross-lingual speech representation learning at scale. *CoRR*, abs/2111.09296.
- Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *CoRR*, abs/2006.11477.
- Theo de Bruin. 2023. Asr-whisper-finetuning. <https://github.com/Theodb/ASR-whisper-finetuning>.
- Tim Dettmers, Mike Lewis, Sam Shleifer, and Luke Zettlemoyer. 2021. 8-bit optimizers via block-wise quantization. *CoRR*, abs/2110.02861.
- Andrea Do, Oscar Brown, Zhengjie Wang, Nikhil Mathew, Zixin Liu, Jawwad Ahmed, and Cheng Yu. 2023. Using fine-tuning and min lookahead beam search to improve whisper. *arXiv:2309.10299*.
- Pelin Dogan-Schönberger, Julian Mäder, and Thomas Hofmann. 2021. *Swissdial: Parallel multidialectal corpus of spoken swiss german*. *Preprint*, arXiv:2103.11401.
- Thomas Palmeira Ferraz, Marcely Zanon Boito, Caroline Brun, and Vassilina Nikoulina. 2023. Multilingual distilwhisper: Efficient distillation of multi-task speech models via language-specific experts. *arXiv:2311.01070*.
- Sanchit Gandhi. 2022. Fine-tune whisper for multilingual asr with transformers. <https://huggingface.co/blog/fine-tune-whisper>.
- Ming-Hao Hsu, Kuan Po Huang, and Hung yi Lee. 2024. Meta-whisper: Speech-based meta-icl for asr on low-resource languages. *Preprint*, arXiv:2409.10429.
- Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Weinberger. 2016. Deep networks with stochastic depth. *Preprint*, arXiv:1603.09382.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526.
- Yunpeng Liu, Xukui Yang, and Dan Qu. 2024. Exploration of whisper fine-tuning strategies for low-resource asr. *EURASIP Journal on Audio, Speech, and Music Processing*, 2024(1):29.
- Rao Ma, Mengjie Qian, Mark J. F. Gales, and Kate M. Knill. 2023. Adapting an asr foundation model for spoken language assessment. *arXiv:2307.09378*.
- OpenAI. 2023. Parameter-efficient fine-tuning of whisper-large v2 in colab on t4 gpu using peft+int8 training. <https://github.com/openai/whisper/discussions/988>.
- Claudio Paonessa, Yanick Schraner, Jan Deriu, Manuela Hürlimann, Manfred Vogel, and Mark Cieliebak. 2023. Dialect transfer for Swiss German speech translation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15240–15254, Singapore. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le. 2019. SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition. In *Proc. Interspeech 2019*, pages 2613–2617.



- Leena G Pillai, Kavya Manohar, Basil K Raju, and Elizabeth Sherly. 2024. [Multistage fine-tuning strategies for automatic speech recognition in low-resource languages](#). *Preprint*, arXiv:2411.04573.
- Andrés Piñeiro-Martín, Carmen García-Mateo, Laura Docio-Fernández, María del Carmen López-Pérez, and Georg Rehm. 2024. [Weighted cross-entropy for low-resource languages in multilingual speech recognition](#). In *Interspeech 2024*, pages 1235–1239. ISCA.
- Michel Plüss, Jan Deriu, Yanick Schraner, Claudio Paonessa, Julia Hartmann, Larissa Schmidt, Christian Scheller, Manuela Hürlimann, Tanja Samardžić, Manfred Vogel, and Mark Cieliebak. 2023. [STT4SG-350: A speech corpus for all Swiss German dialect regions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1763–1772, Toronto, Canada. Association for Computational Linguistics.
- Michel Plüss, Manuela Hürlimann, Marc Cuny, Alla Stöckli, Nikolaos Kapotis, Julia Hartmann, Malgorzata Anna Ulasik, Christian Scheller, Yanick Schraner, Amit Jain, Jan Deriu, Mark Cieliebak, and Manfred Vogel. 2022. [SDS-200: A Swiss German speech to Standard German text corpus](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3250–3256, Marseille, France. European Language Resources Association.
- Michel Plüss, Lukas Neukom, Christian Scheller, and Manfred Vogel. 2021. [Swiss Parliaments Corpus, an Automatically Aligned Swiss German Speech to Standard German Text Corpus](#). In *Swiss Text Analytics Conference 2021*, Proceedings of the Swiss Text Analytics Conference 2021.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Mengjie Qian, Siyuan Tang, Rao Ma, Kate M. Knill, and Mark J.F. Gales. 2024. [Learn and don't forget: Adding a new language to asr foundation models](#). In *Interspeech 2024*, pages 2544–2548.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. [Robust speech recognition via large-scale weak supervision](#). *Preprint*, arXiv:2212.04356.
- Yanick Schraner, Christian Scheller, Michel Plüss, and Manfred Vogel. 2022. [Swiss german speech to text system evaluation](#). *Preprint*, arXiv:2207.00412.
- Aviv Shamsian, Aviv Navon, Neta Glazer, Gill Hetz, and Joseph Keshet. 2024. [Keyword-guided adaptation of automatic speech recognition](#). *arXiv:2406.02649*.
- Clement Sicard, Kajetan Pyszkowski, and Victor Gillioz. 2023. [Spaiche: Extending state-of-the-art asr models to swiss german dialects](#). *Preprint*, arXiv:2304.11075.
- Vaibhav Singh. 2023. [Faster whisper finetuning with lora powered by peft](#). <https://github.com/Vaibhavs10/fast-whisper-finetuning>.
- Nimit S. Sohoni, Christopher R. Aberger, Megan Leszczynski, Jian Zhang, and Christopher Ré. 2022. [Low-memory neural network training: A technical report](#). *Preprint*, arXiv:1904.10631.
- Jakob Straub. 2024. [Die 20 besten schweizer sprichwörter und alltagsweisheiten](#). *Lingoda.com*. Accessed: 2024-12-01.
- Patrick Wilken, Panayota Georgakopoulou, and Evgeny Matusov. 2022. [SubER - a metric for automatic evaluation of subtitle quality](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 1–10, Dublin, Ireland (in-person and online). Association for Computational Linguistics.