

Optimizing Estonian TV Subtitles with Semi-supervised Learning and LLMs

Artem Fedorchenko

Tallinn University of Technology
artem.fedorchenko@taltech.ee

Tanel Alumäe

Tallinn University of Technology
tanel.alumae@taltech.ee

Abstract

This paper presents an approach for generating high-quality, same-language subtitles for Estonian TV content. We fine-tune the Whisper model on human-generated Estonian subtitles and enhance it with iterative pseudo-labeling and large language model (LLM) based post-editing. Our experiments demonstrate notable subtitle quality improvement through pseudo-labeling with an unlabeled dataset. We find that applying LLM-based editing at test time enhances subtitle accuracy, while its use during training does not yield further gains. This approach holds promise for creating subtitle quality close to human standard and could be extended to real-time applications.

1 Introduction

Same-language subtitles for video material, like TV talk shows, investigative pieces, and educational content serve as a valuable resource for the deaf and hard-of-hearing community, non-native speakers and native speakers alike. For instance, recent studies (Mykhalevych and Preply, 2024; Kim et al., 2023) have revealed that 50% of Americans and 85% of the Netflix users overall frequently watch TV and streaming video content with subtitles. Studies show that subtitles can enhance understanding and memory retention. A lot of viewers choose to enjoy their content quietly at home, keeping subtitles on to avoid disturbing their roommates or family.

Subtitles differ from verbatim (word-by-word) transcripts in many aspects. Subtitles represent typically a condensed version of the speech, designed to convey the essential meaning without capturing every word. They may omit filler words, repetitions, and non-verbal sounds, and may rewrite phrases, focusing on clarity and readability for viewers. Since subtitles are displayed on-screen during playback, they are formatted to fit within a limited time frame and limited line length, ensuring they are easy to read while the viewer is watching.

This paper outlines the development of an accurate offline same-language subtitle generation model for Estonian TV content. Using existing human-created subtitles, we fine-tune Whisper (Radford et al., 2022) and explore further improvements with semi-supervised learning and LLM-based post-editing techniques. Our findings demonstrate that Whisper can be trained to closely replicate human subtitling style, creating well-segmented and often rephrased subtitles. Additionally, we find that iterative pseudo-labeling of a large unlabeled dataset improves subtitle quality across all metrics. While a state-of-the-art commercial LLM (OpenAI *gpt-4o*¹) can enhance subtitle quality during test time, its use at training time to improve pseudo-labeled subtitles through post-editing is not effective.

2 Related Work

Both iterative pseudo-labeling and LLM-based post-editing have been an active area of research in the context of verbatim automatic speech recognition (ASR). Pseudo-labeling based semi-supervised learning in ASR has been studied since at least (Zavaliagos et al., 1998) and has been later investigated in several works, e.g. by Vesely et al. (2013); Xu et al. (2020).

To the best of our knowledge, Ma et al. (2023) was the first to show the potential of zero-shot and few-shot LLM-based ASR error correction. This approach has been later extended to take into account uncertainty estimation of ASR outputs (Pu et al., 2023) and retrieval-augmented generation for correcting speech recognition entity name errors (Pusateri et al., 2024).

Xi et al. (2024) showed that LLM-based error correction and data filtering can be also used for refining the pseudo-label transcripts during semi-supervised learning. This work is similar to ours,

¹We used a regular version of GPT-4o, which was accessed on October 14, 2024.

however, it is applied in the context of a code-switched Mandarin-English ASR task.

3 Method

Our method for developing an automated subtitle generation system involves several steps: training with supervised data, using iterative pseudo-labeling, and applying LLM-based error correction.

We start by training the Whisper large-v3 model (Radford et al., 2022) on a supervised dataset. This dataset consists of audio recordings paired with their subtitles.

Next, we use an unsupervised dataset to perform two iterations of pseudo-labeling. In this step, we generate pseudo-labels using the last trained model and combine them with the original supervised dataset, followed by training a new model on this data.

We also apply LLM-based post-editing of the generated subtitles, by instructing the LLM to fix the mistakes in the subtitles and giving it a segment of generated subtitle file. We experiment with applying this LLM-based post-editing in two distinct phases: at test time (i.e., to generated subtitles of the test data) and during training time (i.e., to generated subtitles of the unsupervised dataset).

4 Experiments

4.1 Datasets

As a supervised dataset², we used recordings and the corresponding subtitles from the Estonian national TV. The subtitles had been produced for the deaf and hard-of-hearing community by expert subtitlers. The supervised dataset consists of 993 audio-subtitle pairs, totaling 778 hours of audio, corresponding to 10 different TV show series (multi-party talk shows on various topics, political debates, infotainment programs). We randomly selected 17 recordings out of this set for testing.

The unsupervised dataset contains 7128 audio recordings, amounting to 3923 hours of audio. It contains similar material as the supervised dataset but also contains news program recordings, which the supervised dataset doesn't include.

4.2 Evaluation metrics

While evaluating ASR outputs using word error rate (WER) is relatively straightforward, find-

ing an appropriate metric for evaluating automatic subtitling systems is more complicated. Since subtitles often rephrase spoken content to enhance clarity and readability, WER may not accurately reflect the quality of the subtitles. WER does also not account for the formatting and timing of subtitles, which are crucial for viewer comprehension.

In our work, we use three metrics for comparing machine-generated subtitles against reference subtitles: subtitle edit rate (SubER) (Wilken et al., 2022) and two variations of BLEURT (Selam et al., 2020). SubER is based on a modified version of edit distance that incorporates shifts. This allows it to account for the specific properties of subtitles, such as timing and segmentation. However, SubER doesn't take into account that the same meaning can be conveyed with different words or phrases. Thus, we also use BLEURT for evaluation. BLEURT is a learned metric, trained on subjective human evaluations scores of machine translation references and the corresponding candidate sentences. BLEURT outputs scores that usually in the range of 0..1 (with 1 being a perfect match) and is found to be better correlated with human judgments in several languages than BLEU scores. We used the multilingual BLEURT-20-D12 model introduced by Pu et al. (2021). Furthermore, we use two variations of BLEURT: t-BLEURT and AS-BLEURT, which differ in the way generated subtitles are aligned to references. AS-BLEURT splits the reference subtitles into sentences, aligns generated subtitles to the references (Matusov et al., 2005) and then computes BLEURT score for each sentence, while t-BLEURT does the alignment based on the timing information in the subtitles (Cherry et al., 2021).

4.3 Baseline Model

As a baseline model, we finetuned Whisper on our supervised dataset using a cross-entropy objective. The model was trained for 4 epochs using the AdamW optimizer (Loshchilov and Hutter, 2019) with a learning rate of 1×10^{-5} . We used an effective batch size of 32 audio chunks and applied Stochastic Weight Averaging (SWA) (Izmailov et al., 2018) after the first epoch.

During decoding, we use the Silero VAD model (Silero Team, 2021) to remove non-speech parts.

4.4 Iterative Pseudo-Labeling

Next, to improve performance of the baseline model we used **iterative pseudo-labeling (IPL)**

²<https://cs.taltech.ee/staff/tanel.alumae/data/etv-subtitles/>

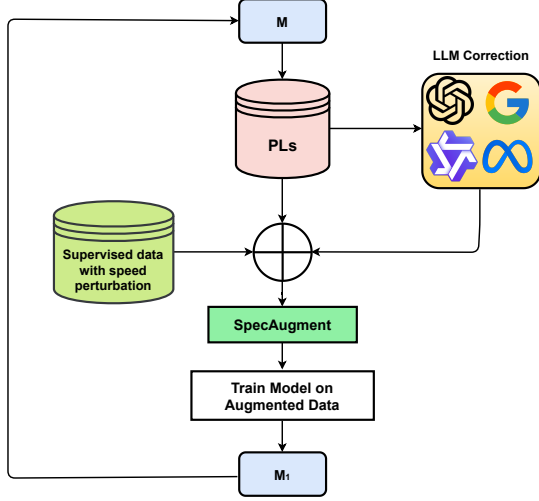


Figure 1: Pseudo-labels generated by model are either passed through LLM or used as is.

— a semi-supervised learning technique that enables model refinement on unlabeled data. Starting with an initial model trained on supervised data, we generate pseudo-labels for unlabeled samples and use these to retrain the model iteratively.

Our approach, which we illustrate in Figure 1, explores two strategies for refining pseudo-labels:

- **Direct Pseudo-Labeling:** Using pseudo-labels generated by the model itself.
- **LLM-Enhanced Pseudo-Labeling:** Refining pseudo-labels with a LLM to correct potential errors and ensure alignment with human subtitling standards.

In both approaches, we combine pseudo-labeled data with the original supervised dataset, modified by applying speed perturbation. To make the model more robust we applied SpecAugment (Park et al., 2019) on spectrogram level.

We did two iterations of training with pseudo-labels, the training setup was similar to the one with supervised data. Additionally, we incorporated weighted loss function:

$$\mathcal{L}_{\text{total}} = (1 - \lambda) \cdot \mathcal{L}_{\text{supervised}} + \lambda \cdot \mathcal{L}_{\text{pseudo-labels}}$$

where $\lambda = 0.35$ was chosen empirically using Optuna (Akiba et al., 2019).

System Instruction:

You are tasked with correcting Estonian subtitles in a subtitle file. **YOU MUST NOT** create, remove, or modify block numbers and timestamps. **ONLY** correct the text within the existing blocks.

Input:

```
1
00:00:00,000 --> 00:00:02,760
Tere õhtust kõigile, algamas
on vestlussaade kahekõne.
2
00:00:02,760 --> 00:00:07,340
Uued rahva poolt palavalt oodatud jõud
on toompeal justkui killustunud.
```

LLM Output:

```
1
00:00:00,000 --> 00:00:02,760
Tere õhtust kõigile, algamas
on vestlussaade "Kahekõne".
2
00:00:02,760 --> 00:00:07,340
Uued rahva poolt palavalt oodatud jõud
on Toompeal justkui killustunud.
```

Figure 2: Example of an LLM instruction used for refining Estonian subtitles. The model corrected the spelling of the TV show name "Kahekõne" and the historical place name "Toompea" in Estonia.

Table 1: Comparison of different LLMs for their performance in error correction.

LLM	SubER↓
-	35.1
GPT-4o	34.2
Llama 3.1 405B (FP8 quant.)	35.5
Qwen 2.5 72B	36.4
Gemma 2 27B	38.4

4.5 LLM-based post-editing

To ensure fast and efficient correction of subtitles using an LLM, we split the generated subtitles into chunks of 40 subtitle blocks. This approach allows for great parallelization without exceeding the maximum token limit per request. An example of the request format is shown in Figure 2.

In the development phase, we evaluated several different LLMs for their suitability for this task. Table 1 shows the SubER results on test data, after applying LLM-based error correction with different LLMs. We compared OpenAI GPT-4o and three of the best open source LLMs from different vendors. As can be seen, only GPT-4o was able to improve SubER-based subtitle accuracy. Based on

Table 2: Results of different models, with or without test-time LLM post-editing.

Finetuning data	Pseudo-label LLM-post-editing?	Test-time LLM-post-editing?	SubER↓	t-BLEURT↑	AS-BLEURT↑
-			59.8	.563	.728
Verbatim transcripts			51.5	.526	.770
Subtitles (A)			35.1	.545	.799
Subtitles (B)		✓	34.2	.582	.810
<i>Pseudo-labeling, iteration 1</i>					
Subtitles + pseudo-labels			34.5	.526	.808
Subtitles + pseudo-labels		✓	33.9	.529	.815
Subtitles + pseudo-labels	✓		34.4	.525	.810
Subtitles + pseudo-labels	✓	✓	33.9	.528	.816
<i>Pseudo-labeling, iteration 2</i>					
Subtitles + pseudo-labels			33.4	.529	.853
Subtitles + pseudo-labels (C)		✓	33.1	.598	.858
Subtitles + pseudo-labels	✓		33.6	.570	.854
Subtitles + pseudo-labels	✓	✓	33.3	.571	.856

these results, we used GPT-4o in our experiments.

During our experiments, we observed that LLMs often struggle to output the exact timestamps and block numbers correctly. To address this, we verified these details against the original subtitles to ensure accuracy and re-requested the LLM to fix the issue, if necessary. We also experimented with one-shot and few-shot prompts but did not observe any significant quality improvement, so we opted not to include them. Additionally, we set a threshold on the number of allowable reference check failures: if the model failed more than 3 times, we reverted to the original subtitle.

4.6 Results

Table 2 lists evaluation results of the native Whisper model (not fine-tuned on additional data), Whisper fine-tuned on 1066 hours of verbatim transcripts from the TalTech Estonian Speech Dataset 1.0 (Alumäe et al., 2023), and after fine-tuning with different sets of subtitle datasets. The table also highlights the effects of LLM-based post-editing applied during both the training and testing phases.

The results indicate that fine-tuning on subtitle data yields notably lower SubER values compared to fine-tuning on verbatim transcripts, demonstrating the different nature of subtitles and verbatim transcripts. However, the BLEURT scores for both the native Whisper model and the version fine-tuned on verbatim transcripts are surprisingly high. This outcome may be attributed to BLEURT’s design as a semantic similarity metric,

which effectively maps both verbatim transcripts and subtitle-like compressed transcripts to proximate points in its semantic space.

To support our interpretation of the achieved results, we computed Wilcoxon signed-rank test (Wilcoxon, 1945) between models **A**, **B** and **C** highlighted in the Table 2. P-value achieved from comparing model **A** to **B** is 0.000, **B** to **C** is 0.004 and **A** to **C** is 0.000. These p-values are all below common significance thresholds (e.g., 0.05), indicating that the differences between the models are statistically significant.

Given that, findings suggest that iterative semi-supervised learning enhances subtitle quality, as evidenced by improvements across all test metrics. LLM-based post-editing applied to decoded subtitles provides additional benefits in most cases. However, contrary to findings in (Xi et al., 2024), applying LLM-based post-editing to pseudo-labeled subtitles in the unsupervised dataset does not yield further improvements.

Although a formal human evaluation of the generated subtitles was not conducted, the authors’ subjective assessment suggests that minimal manual post-editing would be required to achieve error-free subtitles, particularly for in-domain TV data. A sample video from our test dataset, featuring both reference subtitles and subtitles generated by our best model³ is available at <https://www.youtube.com/watch?v=bEow5vGIgZc>. A smaller version of

³<https://huggingface.co/TalTechNLP/whisper-large-v3-et-sub>

this model based on Whisper *large-v3-turbo* can be freely used via a simple web application⁴.

5 Conclusion

In this work, we presented an approach to automated subtitle generation, leveraging the multilingual Whisper model, semi-supervised learning, and LLM-based post-editing. By utilizing supervised and unsupervised datasets, we demonstrated that iterative pseudo-labeling can indeed improve the quality of subtitles. Our results show that applying an LLM during test time has a more significant impact on the results across all the key metrics than during training time. Future work will focus on adapting our approach to real-time scenarios.

References

- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. Op-tuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Tanel Alumäe, Joonas Kalda, Külliki Bode, and Martin Kaitsa. 2023. Automatic closed captioning for Estonian live broadcasts. In *Proc. NoDaLiDa*, Tórshavn, Faroe Islands.
- Colin Cherry, Naveen Arivazhagan, Dirk Padfield, and Maxim Krikun. 2021. Subtitle translation as markup translation. In *Proc. Interspeech*, pages 2237–2241.
- Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. 2018. Averaging weights leads to wider optima and better generalization. In *Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 876–885. AUAI Press.
- Hyunju Kim, Yan Tao, Chuanrui Liu, Yuzhuo Zhang, and Yuxin Li. 2023. Comparing the impact of professional and automatic closed captions on video-watching experience. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*.
- Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Rao Ma, Mengjie Qian, Potsawee Manakul, Mark Gales, and Kate Knill. 2023. Can generative large language models perform asr error correction? *arXiv preprint arXiv:2307.04172*.
- Evgeny Matusov, Gregor Leusch, Oliver Bender, and Hermann Ney. 2005. Evaluating machine translation output with automatic sentence segmentation. In *Proc. IWSLT*.
- Nadiia Mykhalevych and Preply. 2024. Survey: Why America is obsessed with subtitles. Accessed: 2024-10-25.
- 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 *Interspeech 2019*. ISCA.
- Amy Pu, Hyung Won Chung, Ankur P Parikh, Sebastian Gehrmann, and Thibault Sellam. 2021. Learning compact metrics for MT. In *Proc. EMNLP*.
- Jie Pu, Thai-Son Nguyen, and Sebastian Stüker. 2023. Multi-stage large language model correction for speech recognition. *ArXiv*, abs/2310.11532.
- Ernest Pusateri, Anmol Walia, Anirudh Kashi, Bortik Bandyopadhyay, Nadia Hyder, Sayantan Mahinder, R. Anantha, Daben Liu, and Sashank Gondala. 2024. Retrieval augmented correction of named entity speech recognition errors. *ArXiv*, abs/2409.06062.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. Robust speech recognition via large-scale supervision. *arXiv preprint arXiv:2212.04356*.
- Thibault Sellam, Dipanjan Das, and Ankur P Parikh. 2020. BLEURT: Learning robust metrics for text generation. In *Proc. ACL*.
- Silero Team. 2021. Silero vad: Pre-trained enterprise-grade voice activity detector. <https://github.com/snakers4/silero-vad>.
- Karel Veselý, Mirko Hannemann, and Lukáš Burget. 2013. Semi-supervised training of deep neural networks. In *2013 IEEE Workshop on Automatic Speech Recognition and Understanding*, pages 267–272. IEEE.
- Frank Wilcoxon. 1945. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83.
- Patrick Wilken, Panayota Georgakopoulou, and Evgeny Matusov. 2022. SubER: A metric for automatic evaluation of subtitle quality. *arXiv preprint arXiv:2205.05805*.
- Yu Xi, Wen Ding, Kai Yu, and Junjie Lai. 2024. Semi-supervised learning for code-switching ASR with large language model filter. *arXiv preprint arXiv:2407.04219*.
- Qiantong Xu, Tatiana Likhomanenko, Jacob Kahn, Awni Hannun, Gabriel Synnaeve, and Ronan Collobert. 2020. Iterative pseudo-labeling for speech recognition. *arXiv preprint arXiv:2005.09267*.

⁴<https://huggingface.co/spaces/TanelAlumae/whisper-large-v3-et-subs>

George Zavaliagkos, Man-Hung Siu, Thomas Colthurst, and Jayadev Billa. 1998. Using untranscribed training data to improve performance. In *Proc. ICSLP*, volume 1998.