

Swiss Parliaments Corpus, an Automatically Aligned Swiss German Speech to Standard German Text Corpus

Michel Plüss Lukas Neukom Christian Scheller Manfred Vogel

Institute for Data Science

University of Applied Sciences and Arts Northwestern Switzerland

Windisch, Switzerland

michel.pluess@fhnw.ch

Abstract

We present the Swiss Parliaments Corpus (SPC), an automatically aligned Swiss German speech to Standard German text corpus. This first version of the corpus is based on publicly available data of the Bernese cantonal parliament and consists of 293 hours of data. It was created using a novel forced sentence alignment procedure and an alignment quality estimator, which can be used to trade off corpus size and quality. We trained Automatic Speech Recognition (ASR) models as baselines on different subsets of the data and achieved a Word Error Rate (WER) of 0.278 and a BLEU score of 0.586 on the SPC test set. The corpus is freely available for download¹.

1 Introduction

Swiss German is a family of dialects spoken by around five million people in Switzerland. It is different from Standard German regarding phonetics, vocabulary, morphology, and syntax. Swiss German is mostly a spoken language. While it is also used in writing, particularly in informal text messages, it lacks a standardized writing system. This leads to difficulties for automated text processing such as spelling ambiguities and a huge vocabulary size. For Swiss German ASR, we therefore focus on end-to-end approaches from Swiss German speech to Standard German text. This can be viewed as a Speech Translation problem with similar source and target languages. For example,

the Swiss German sentence "Ide Abfahrt hetter de sächsti Platz beleit" can be translated to the Standard German sentence "In der Abfahrt belegte er den sechsten Platz". Here, the past tense changes.

Currently, training an ASR model for Swiss German is challenging due to the lack of public training data. Only a few hours of Swiss German speech with Standard German text are available. To reach high-quality ASR results, a corpus with thousands of hours of transcribed speech is required. For example, Park et al. (2020) set the current state-of-the-art on the English LibriSpeech (Panayotov et al., 2015) test-other benchmark with a WER of 0.034 using 960 hours of labeled training data and another 57700 hours of unlabeled data.

While there is no ready-to-use training data, many Swiss parliaments record their debates. Most communal and some cantonal parliaments hold their meetings in Swiss German. Some of them do a full transcript of the recordings in Standard German resulting in more than 4000 hours of raw data.

To transform the raw data into training data, we developed a novel forced sentence alignment algorithm which handles the problems created by the language mismatch between audio and text such as changes to the word order within a sentence (sentence reordering). It is based on a German ASR model and global alignment and includes a learned filter component specifically tuned for the Swiss German speech to Standard German text use case². Using the developed alignment algorithm, we created and published a first corpus called the Swiss Parliaments Corpus consisting of data from the parliament Grosser Rat Kanton Bern.

The remainder of this paper is structured as follows: Related work is discussed in section 2. The

Copyright © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0)

¹<https://www.cs.technik.fhnw.ch/i4ds-datasets>

²The code is available in our GitHub repository: <https://github.com/festivalhopper/swiss-parliaments-corpus-paper>

forced sentence alignment procedure is described in section 3. Details about our corpus can be found in section 4. Section 5 contains baseline models and experiments. Section 6 wraps up the paper and gives directions for future work.

2 Related Work

An earlier version of this corpus was previously published as part of the Low-Resource Speech-to-Text shared task at GermEval 2020 (GermEval Task 4) (Plüss et al., 2020). To our knowledge, there are only two other publicly available corpora for Swiss German ASR. ArchiMob (Samardžić et al., 2016) includes 69 hours of Swiss German speech and corresponding Swiss German transcripts. Unfortunately, Standard German transcripts are not available. The Radio Rottu Oberwallis dataset (Garner et al., 2014) includes 8 hours of speech, of which only 2 hours have Standard German transcripts in addition to Swiss German transcripts. Furthermore, the Standard German dataset of the Common Voice (Ardila et al., 2019) ASR corpus has 1 % of its utterances spoken in a Swiss German accent, which however strongly differs from actual Swiss German speech.

There are different approaches to forced alignment of long speech recordings in the context of creating an ASR corpus. Our procedure is similar to that described in Hazen (2006); Panayotov et al. (2015); Pratap et al. (2020). Like our approach, these methods initially transcribe audios using an ASR system, followed by an alignment stage and a final refinement stage. The main differences are that these approaches do not yield strictly sentence-level alignments, which are a requirement for our work due to the possibility of sentence reorderings between Standard German and Swiss German, and our novel approach to filter the corpus and improve the quality.

3 Forced Sentence Alignment Procedure

Our forced sentence alignment procedure takes a Swiss German recording of arbitrary length and the corresponding manual Standard German transcript as inputs. The audio file is transcribed with an ASR model. An important requirement for this model is the ability to annotate accurate start and end times of each word in the output. Since no publicly available Swiss German ASR model with this feature exists, we resort to a Standard German model. The ASR transcript is then globally aligned

to the manual transcript using the Biopython (Cock et al., 2009) implementation of the Needleman-Wunsch algorithm (Needleman and Wunsch, 1970). The manual transcript is split into sentences using spaCy (Honnibal and Montani, 2017). Each of these sentences is mapped to a start and end time in the recording via the global alignment and the per-word start and end times provided by the ASR model.

3.1 Alignment Corpus and Metrics

We created a separate internal alignment corpus to be able to measure the quality of our sentence alignment. It consists of almost 6 hours of transcribed recordings from four different parliaments and one other data source. We split the corpus into a training and a test set (60-40 split). Recordings and transcripts were manually sentence-aligned.

We define an aligned sentence as a three-tuple of the sentence, its start and end time. We call an aligned sentence empty if the start and end times are not set, which means the sentence is not spoken in the recording. This happens because transcripts sometimes have errors such as missing or additional sentences.

Our goal is to maximize the following metrics during the creation of the corpus:

- **The Intersection over Union (IoU)** reflects the alignment quality. We report the mean IoU over all predicted aligned sentences for which the manual as well as the predicted aligned sentence are not empty.
- **The sentence precision and recall** reflect the corpus quality and size. A predicted aligned sentence counts as true positive (TP) if the manual as well as the predicted aligned sentence are not empty. True negative (TN) means the manual as well as the predicted aligned sentence are empty. False positive (FP) means the manual aligned sentence is empty, but the predicted aligned sentence is not empty. False negative (FN) means the manual aligned sentence is not empty, but the predicted aligned sentence is empty. The sentence precision is equal to $TP / (TP + FP)$. The sentence recall is equal to $TP / (TP + FN)$.

3.2 IoU Estimate Filter and Further Refinements

We filter out sentences with a bad alignment quality based on an estimate of their IoU. We fit a Gradient

Hyperparameter	Value
num_leaves	3
min_child_samples	7
max_bin	7597

Table 1: LightGBM hyperparameters for the IoU regressor. For all other parameters we use the default values as defined by the LightGBM authors.

Boosting regressor to estimate a sentence’s IoU using the following features:

- **Length ratio** of the manual transcript sentence to the part of the ASR transcript it was aligned to
- **Alignment score** of the manual transcript sentence, normalized by its length
- **Mean speech recognition confidence** as reported by the ASR system over the words the manual transcript sentence was aligned to
- **Chars per second**, i.e. ratio of the manual transcript sentence length to the audio length (predicted aligned sentence end time minus start time)

We use the LightGBM implementation by Ke et al. (2017). Table 1 shows our hyperparameters. These were found using Bayesian optimization.

The regressor estimates the IoU in a 3-fold cross validation experiment on the training set of our alignment corpus with a mean absolute error of 0.108 (IoU values are in the interval [0, 1]). We propose two different IoU estimate thresholds. A threshold of 0.7 is supposed to keep as many sentences as possible and only discard sentences with a bad alignment quality, e.g. for a training set. A threshold of 0.9 is supposed to keep only sentences with a very good alignment quality, e.g. for a test set. Thresholds were found using a parameter sweep on the test set of the alignment corpus.

Two more refinements were implemented: to filter out manual transcripts that are clearly mismatched or incomplete, no alignment is created if the length ratio of the longer transcript to the shorter transcript is greater than six. The optimal ratio was found using a parameter sweep on the alignment corpus test set. Finally, we fit a start and end time correction offset on the training set of our alignment corpus and calibrate the start and

ASR Model	WER
Amazon Transcribe	0.626
Google Speech-to-Text	0.725

Table 2: Comparison of the performance of different ASR models on the GermEval 2020 Task 4 public test set

end times of each sentence by adding the correction offset. This leads to a minor IoU improvement because the times reported by the ASR model can be slightly off.

3.3 Experiments and Results

We evaluated two ASR models, Amazon Transcribe³⁴ and Google Speech-to-Text⁵⁶, on the public test set of GermEval 2020 Task 4 (Plüss et al., 2020). Table 2 shows the results of this comparison. Amazon is ahead of Google by 0.1 WER. This suggests that Amazon succeeded in improving the performance for Swiss German with its specialized model, but still has a long way to go to achieve a general-purpose ASR model with a WER comparable to English or Standard German models. In comparison to the performance of the winning contribution to GermEval 2020 Task 4 by Büchi et al. (2020) with a WER of 0.403, Amazon Transcribe is more than 0.2 WER behind.

We conducted experiments using Amazon Transcribe and Google Speech-to-Text as the ASR engines with different combinations of refinements⁷. We determined the parameters for the global alignment algorithm using Bayesian optimization with 3-fold cross validation on the training set of our alignment corpus. They are listed in appendix A. Table 3 shows the results. The lead in WER by 0.1 (see table 2) for Amazon translates to a mean IoU that is 0.151 higher than Google’s result with the same settings. The 0.045 advantage of the latter in sentence recall does not make up for this, even less so because we prefer quality over quantity. For

³<https://aws.amazon.com/transcribe>

⁴We used the "Swiss German" model. Based on the information we found, we believe this is a model for Standard German, specialized on Swiss accents, not for actual Swiss German.

⁵<https://cloud.google.com/speech-to-text>

⁶We used the "German (Germany)" model, "German (Switzerland)" was not yet available at the time of the experiment.

⁷We could not use the Büchi et al. model because it does not provide word start and end times

ASR Model	Settings	Mean IoU	Sentence Precision	Sentence Recall
Amazon Transcribe	No Refinements	0.832	1.000	0.956
Amazon Transcribe	Length Ratio	0.836	1.000	0.949
Amazon Transcribe	Time Calibration	0.836	1.000	0.956
Amazon Transcribe	Length Ratio + Time Calibration	0.840	1.000	0.949
Google Speech-to-Text	Length Ratio + Time Calibration	0.689	1.000	0.994
Amazon Transcribe	Length Ratio + Time Calibration	0.888	1.000	0.822
	+ IoU Estimate Filter 0.7			
Amazon Transcribe	Length Ratio + Time Calibration	0.927	1.000	0.488
	+ IoU Estimate Filter 0.9			

Table 3: Sentence alignment metrics on the test set of our alignment corpus for two ASR models with various settings

Amazon Transcribe, enabling length ratio filtering as well as time calibration appears to be the best option, resulting in a mean IoU of 0.840 and a sentence recall of 0.949. The alignment quality can be further improved using the IoU estimate filter. A threshold of 0.7 leads to an increase of 0.048 in mean IoU and a decrease of 0.127 in sentence recall, whereas a threshold of 0.9 leads to an increase of 0.087 in mean IoU and a decrease of 0.461 in sentence recall. The sentence precision is perfect in all experiments.

4 Swiss Parliaments Corpus

Using our forced sentence alignment procedure, we created and published⁸ a corpus called the Swiss Parliaments Corpus. It is based on recordings and transcripts from the parliament Grosser Rat Kanton Bern⁹. As expected, given the location of the parliament, most speakers have a Bernese dialect. The recordings are MP4 videos, one video per parliament meeting, with a length spanning from 28 minutes to 4 hours and 2 minutes. The transcripts are in PDF format, with one PDF containing a whole session with usually around 10 to 15 meetings.

4.1 Corpus Parts

Table 4 gives an overview of the different corpus parts and their sizes. We created an unfiltered training set called `train_all` with 293 hours of data. We then used IoU estimate filtering to create two training subsets, `train_0.7` with a threshold of 0.7 and

⁸<https://www.cs.technik.fhnw.ch/i4ds-datasets>

⁹<https://www.gr.be.ch/gr/de/index/sessionen/sessionen.html>

Corpus Part	Audio Length in Hours	Number of Speakers
Raw data	460	-
<code>train_all</code>	293	198
<code>train_0.7</code>	256	195
<code>train_0.9</code>	176	194
<code>test</code>	6	26

Table 4: Overview of the different subsets of the corpus, their sizes and the number of unique speakers

256 hours of data as well as `train_0.9` with a threshold of 0.9 and 176 hours of data. The unfiltered training set contains an IoU estimate column to create a training set with a custom threshold. The test set was created with a threshold of 0.9 and contains 6 hours of data. We could therefore transform 65 % of the raw data to training or test data.

4.2 Settings, Filters, Split

We used semi-global alignment parameters (no gap penalties on the start and end of both sequences, see appendix B) to deal with incomplete recordings and additional irrelevant recorded audio. Length ratio filtering was disabled while time calibration was enabled. We applied the following additional filters:

- Chars per second must be between 6 and 23. The average for chars per second in this corpus is 15. Aligned sentences outside of this range probably either contain a lot of idle time in the recording or additional text that is not recorded.

- We detect the language of each sentence using `langdetect`¹⁰ and only keep German sentences.
- (Test set only) Audio length must be at least 1 second.
- (Test set only) Audio length must be less than 15 seconds.
- (Test set only) Sentences must be unique across the whole dataset.

Speakers are automatically deduplicated. The train-test split guarantees that the utterances of a speaker are only contained in either the training set or the test set, never in both. To ensure that the speakers in the test set are diverse enough, a speaker can only be part of the test set if her or his utterances make up less than 10 % of the whole test set.

5 ASR Baselines

All the baseline models are implemented using the ESPnet framework (Watanabe et al., 2018). We trained a Transformer model (Vaswani et al., 2017) as well as a Conformer model (Gulati et al., 2020). The network architectures closely follow the Common Voice example of ESPnet¹¹ using a hybrid CTC/attention encoder-decoder framework (Watanabe et al., 2017).

Inputs are first down-sampled to 1/4 length by two strided 2D convolution layers and ReLU activations. The Transformer encoder consists of 12 self-attention blocks with 2048 units. Similarly, the Conformer encoder uses 12 Conformer-layers with 2048 units. Both Transformer and Conformer models use a Transformer decoder with six self-attention blocks with 2048 units.

For the training of both models we use the Adam optimizer (Kingma and Ba, 2015) and a warmup learning rate schedule similar to the one proposed in Vaswani et al. (2017), but with a fixed warmup period of 25000 steps and a maximum learning rate of 0.002. As input we use 80-channel log-mel filterbanks that are shifted to have a mean of zero. Speed perturbation (Ko et al., 2015) with random factors between 0.9 and 1.1 and SpecAugment (Park et al., 2019) are used for data augmentation. Both models are trained with a combined Connectionist Temporal Classification (CTC) (Graves et al., 2006)

¹⁰<https://pypi.org/project/langdetect>

¹¹<https://github.com/espnet/espnet/tree/master/egs2/commonvoice/asr1>

Model	Dataset	WER	BLEU
Transformer	train_all	0.297	0.548
	train_0.7	0.293	0.553
	train_0.9	0.303	0.537
Conformer	train_all	0.289	0.577
	train_0.7	0.278	0.586
	train_0.9	0.287	0.577

Table 5: Test WER and BLEU scores of Transformer and Conformer models on all subsets of the SPC.

and cross entropy loss with weights 0.3 and 0.7, respectively.

During decoding we use a 16-layer Transformer language model, trained on the SPC texts as well as the German texts of the EuroParl Corpus (Koehn, 2005), using beam search with a beam size of 50.

We trained both models on all subsets of the SPC until convergence for 200 epochs. The results of all models are shown in Table 5. We report WER, commonly used to evaluate ASR systems, as well as the BLEU (Papineni et al., 2002) score, commonly used to evaluate Machine Translation and Speech Translation systems. For BLEU, we use the implementation provided by NLTK (Bird et al., 2009) with default parameters. In our experiments, WER and BLEU show a negative correlation as expected (WER: lower is better, BLEU: higher is better), indicating that both are similarly useful metrics.

The Conformer model performs better than the Transformer model in all experiments. This is in line with the findings of Gulati et al. (2020) on the LibriSpeech benchmark. For both models, all dataset splits resulted in similar WER and BLEU scores, even though their sizes differ significantly (see Table 4). This can be explained by the higher quality when filtering based on IoU scores. As a result, we can train models on 60 % of the data, resulting in faster training, without losing performance.

6 Conclusion

In this work, we introduced the Swiss Parliament Corpus, an automatically aligned Swiss German speech to Standard German text corpus. We proposed a multi-stage forced sentence alignment procedure that leverages existing Standard German ASR systems and uses a novel IoU estimator for refinement. We also provided Transformer and

Conformer ASR baseline models that showcase the benefits of the IoU estimates. The best model achieves a WER of 0.278 and a BLEU score of 0.586 on the SPC test set.

We believe that our forced sentence alignment procedure is a step towards making large-vocabulary speech recognition for all Swiss German dialects possible. The SPC with its 293 hours of training data supports this thesis. It is freely available for download¹².

In future work, we plan to increase the corpus size and the dialect diversity by aligning recordings and transcripts of additional parliaments. We also plan to collect data for a test set representing all Swiss German dialects since the SPC is domain-specific and includes mostly Bernese speakers. This would facilitate a fair comparison for Swiss German ASR systems. In the future, we will also use the trained models to improve the forced sentence alignment algorithm results in a similar way as [Sennrich and Volk \(2011\)](#).

Furthermore, the effects of the IoU filter on the quality of the dataset as well as the impact on the model need further investigation. Finally, we want to investigate the correlation of WER and BLEU with human evaluation to understand which metric is most appropriate for the problem. In this context, it would also be interesting to further investigate the differences between a literal transcription in Swiss German and a Standard German translation.

Acknowledgments

First and foremost, we would like to thank the parliamentary services of the canton of Bern for their work on the transcription of the debates and for publishing recordings and transcripts on their website. Without them, the SPC would not exist.

Furthermore, we thank Pascal Thormeier who created the first version of our alignment corpus during his bachelor’s thesis.

We also thank the participants of GermEval 2020 Task 4 for the fruitful discussions in the aftermath of the task, which lead to several improvements of the SPC.

References

Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben

¹²<https://www.cs.technik.fhnw.ch/i4ds-datasets>

Morais, Lindsay Saunders, Francis M Tyers, and Gregor Weber. 2019. Common voice: A massively-multilingual speech corpus. *arXiv preprint arXiv:1912.06670*.

Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. ” O’Reilly Media, Inc.”.

Matthias Büchi, Malgorzata Anna Ulasik, Manuela Hürlimann, Fernando Benites, Pius von Däniken, and Mark Cieliebak. 2020. Zhaw-init at germeval 2020 task 4: Low-resource speech-to-text. In *SWISSTEXT & KONVENS 2020*, Proceedings of the 5th Swiss Text Analytics Conference (SwissText) & 16th Conference on Natural Language Processing (KONVENS).

Peter J. A. Cock, Tiago Antao, Jeffrey T. Chang, Brad A. Chapman, Cymon J. Cox, Andrew Dalke, Iddo Friedberg, Thomas Hamelryck, Frank Kauff, Bartek Wilczynski, and Michiel J. L. de Hoon. 2009. Biopython: freely available Python tools for computational molecular biology and bioinformatics. *Bioinformatics*, 25(11):1422–1423.

Philip Garner, David Imseng, and Thomas Meyer. 2014. Automatic speech recognition and translation of a swiss german dialect: Walliserdeutsch. In *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*.

Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pages 369–376.

Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. 2020. Conformer: Convolution-augmented Transformer for Speech Recognition. In *Proceedings of Interspeech*, pages 5036–5040.

Timothy J. Hazen. 2006. Automatic alignment and error correction of human generated transcripts for long speech recordings. In *Proceedings of Interspeech*.

Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.

Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. Lightgbm: A highly efficient gradient boosting decision tree. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, page 3149–3157.

- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015*.
- Tom Ko, Vijayaditya Peddinti, Daniel Povey, and Sanjeev Khudanpur. 2015. Audio augmentation for speech recognition. In *Sixteenth Annual Conference of the International Speech Communication Association*.
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *MT summit*, volume 5, pages 79–86.
- Saul B. Needleman and Christian D. Wunsch. 1970. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *J. Mol. Biol.*, 48:443–453.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: An ASR corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP*, pages 5206–5210.
- 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.
- 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. *arXiv preprint arXiv:1904.08779*.
- Daniel S. Park, Yu Zhang, Ye Jia, Wei Han, Chung-Cheng Chiu, Bo Li, Yonghui Wu, and Quoc V Le. 2020. Improved noisy student training for automatic speech recognition. In *Proceedings of Interspeech*.
- Michel Plüss, Lukas Neukom, and Manfred Vogel. 2020. Germeval 2020 task 4: Low-resource speech-to-text. In *SWISSTEXT & KONVENS 2020*, Proceedings of the 5th Swiss Text Analytics Conference (SwissText) & 16th Conference on Natural Language Processing (KONVENS).
- Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert. 2020. Mls: A large-scale multilingual dataset for speech research. *arXiv preprint arXiv:2012.03411*.
- Tanja Samardžić, Yves Scherrer, and Elvira Glaser. 2016. Archimob - a corpus of spoken swiss german. Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016), pages 4061–4066.
- Rico Sennrich and Martin Volk. 2011. [Iterative, MT-based sentence alignment of parallel texts](#). In *Proceedings of the 18th Nordic Conference of Computational Linguistics (NODALIDA 2011)*, pages 175–182, Riga, Latvia. Northern European Association for Language Technology (NEALT).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30.
- Shinji Watanabe, Takaaki Hori, Shigeki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplín, Jahn Heymann, Matthew Wiesner, Nanxin Chen, Adithya Renduchintala, and Tsubasa Ochiai. 2018. ESPnet: End-to-end speech processing toolkit. In *Proceedings of Interspeech*, pages 2207–2211.
- Shinji Watanabe, Takaaki Hori, Suyoun Kim, John R Hershey, and Tomoki Hayashi. 2017. Hybrid ctc/attention architecture for end-to-end speech recognition. *IEEE Journal of Selected Topics in Signal Processing*, 11(8):1240–1253.

A Alignment Parameters Optimized on Alignment Corpus

Alignment Parameter	Value
match_score	0.039
mismatch_score	−1.000
truth_left_open_gap_score	−0.504
truth_internal_open_gap_score	−1.000
truth_right_open_gap_score	−0.440
truth_left_extend_gap_score	−0.244
truth_internal_extend_gap_score	−0.482
truth_right_extend_gap_score	−0.259
stt_left_open_gap_score	−1.000
stt_internal_open_gap_score	−0.770
stt_right_open_gap_score	−0.982
stt_left_extend_gap_score	−0.253
stt_internal_extend_gap_score	−0.770
stt_right_extend_gap_score	−0.562

Table 6: Alignment parameters found using Bayesian optimization with 3-fold cross validation on the training set of our alignment corpus

B Alignment Parameters for SPC

Alignment Parameter	Value
match_score	1.0
mismatch_score	-1.0
truth_left_open_gap_score	0.0
truth_internal_open_gap_score	-1.0
truth_right_open_gap_score	0.0
truth_left_extend_gap_score	0.0
truth_internal_extend_gap_score	-1.0
truth_right_extend_gap_score	0.0
stt_left_open_gap_score	0.0
stt_internal_open_gap_score	-1.0
stt_right_open_gap_score	0.0
stt_left_extend_gap_score	0.0
stt_internal_extend_gap_score	-1.0
stt_right_extend_gap_score	0.0

Table 7: Alignment parameters used to create the SPC