

Dialectal Speech Recognition and Translation of Swiss German Speech to Standard German Text: Microsoft’s Submission to SwissText 2021

Yuriy Arabskyy, Aashish Agarwal, Subhadeep Dey, Oscar Koller

Microsoft - Munich, Germany

{yuarabsk, t-aagarwal, subde, oskoller} @microsoft.com

Abstract

This paper describes the winning approach in the Shared Task 3 at SwissText 2021 on Swiss German Speech to Standard German Text, a public competition on dialect recognition and translation. Swiss German refers to the multitude of Alemannic dialects spoken in the German-speaking parts of Switzerland. Swiss German differs significantly from standard German in pronunciation, word inventory and grammar. It is mostly incomprehensible to native German speakers. Moreover, it lacks a standardized written script. To solve the challenging task, we propose a hybrid automatic speech recognition system with a lexicon that incorporates translations, a 1st pass language model that deals with Swiss German particularities, a transfer-learned acoustic model and a strong neural language model for 2nd pass rescoring. Our submission reaches 46.04% BLEU on a blind conversational test set and outperforms the second best competitor by a 12% relative margin.

1 Introduction

While general speech recognition has matured to a point where it surpasses human performance on specific datasets (Xiong et al., 2016), dialectal recognition as in the case of Swiss German (Nigmatulina et al., 2020) or Arabic dialects (Ali et al., 2021; Hussein et al., 2021; Ali, 2018) still represents a major challenge. Swiss German refers to the multitude of Alemannic dialects spoken in the German-speaking parts of Switzerland. It hence

represents dialects that differ significantly from standard or high German in pronunciation, word inventory and grammar. Moreover, it lacks a standardized writing system. High German is commonly used for the large majority of written communication between people and in the media of German-speaking Switzerland, while in rather informal chats and short messaging Swiss people may use transliterated non-standardized Swiss dialect. The process of transcribing Swiss German into high German text therefore requires speech recognition with an inherent translation step. Moreover, the task can be considered low-resource as available data remains extremely scarce.

In previous studies, Garner et al. (2014) tackled this challenge by training hybrid models (HMM-GMM, HMM-DNN, and KL-HMM) to transcribe Walliserdeutsch, a Swiss-German dialect spoken in the south-western alpine canton of Switzerland and further used a phrase-based machine translation model to translate it to standard German. Following this, other researchers explored techniques to add the translation step in the lexicon by directly mapping Swiss-German pronunciation to standard German. Stadtschnitzer and Schmidt (2018) estimated Swiss German pronunciations from a standard German speech recognition model using a data-driven technique, and trained stronger TDNN-LSTM based acoustic models. Kew et al. (2020) and Nigmatulina et al. (2020) trained transformer-based G2P models from standard German to Swiss pronunciations and trained a Kaldi-based TDNN+ivector system using the WSJ recipe¹. Yet a third approach is to directly apply end-to-end deep learning models. Büchi et al. (2020) and Agarwal and Zesch (2020) at SwissText 2020 used the Jasper architecture (Li et al., 2019) and Mozilla DeepSpeech (Hannun et al.,

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

¹https://github.com/kaldi-asr/kaldi/tree/master/egs/ws_j

2014), respectively. In both cases the system was first trained on high German data and then transferred to Swiss German.

In this paper, we describe our proposal to solve the challenging task of transcribing Swiss German speech to standard German text. It won the competition at SwissText 2021 with a large margin to other competing systems.

2 System Overview

In this section, we propose our changes to a conventional hybrid (Bouillard and Dupont, 1996) automatic speech recognition (ASR) system, which relies on lexicon and alignments for good performance. We present details in order to enable it for dialect speech recognition and translation.

2.1 Data

To train our proposed model, we utilized a selection of publicly available and internal datasets. Our starting point was the Swiss Parliament Corpus V2 dataset (Plüss et al., 2020) shared as part of the SwissText 2021 competition. It covers 293 hours and contains recordings from the local parliament of the Kanton Bern. Its transcripts are in standard German while the audio covers Swiss German (predominantly in the Bernese dialect). The dataset has been preprocessed by the publishers with the purpose of cleaning its annotations and ensuring a good match between audio content and transcription. It is provided with a choice of different preprocessing flavors. We used the `train_all` split. In addition, we used a 493-hour internal dataset representing a media domain encompassing conversational speech from interviews, discussions, podcasts and others. A subset (around 50 hours) of the data is annotated with both Swiss transliterations as well as standard German. The remaining data has only been annotated with standard German. Additionally, we used an internal high German dataset encompassing around 10k hours to pre-train our model.

In terms of test data, the SwissText 2021 competition was accompanied by a 13 hours conversational test set covering Swiss German speakers from all German-speaking parts of Switzerland. The encountered dialectal distribution is claimed to closely match the real distribution in Switzerland. The set was not disclosed to the participants. Hence, for the analysis in this paper, we report our numbers on a publicly available test set part of the

dataset from the Bernese parliament (Plüss et al., 2020). It comprises 6 hours of dialectal speech.

2.2 Lexicon

We propose to incorporate the translation from Swiss German to standard German as part of the lexicon. However, this leads to a complex and often ambiguous mapping from phoneme to grapheme sequences, which is very different from languages with a direct relation between writing scheme and pronunciation (e.g. English or standard German). Subsequently, statistical models that map graphemes to phonemes (G2P) trained on Swiss German data incorporating such translations yield much noisier output with significantly higher phone error rates as compared to G2Ps for standard languages. To mitigate this problem, we construct the lexicon in several stages.

In a first step, we make use of parallel corpora encompassing Swiss and standard German annotations to extract word mappings between Swiss and standard German. Sophisticated filtering methods help to ensure a high quality of these mappings. We opt for frequency filtering and filtering based on vicinity in a word embedding space (Bojanowski et al., 2017) of Swiss German words taking the most frequent mapping as center point.

In a second step, a standard German G2P model is applied to convert Swiss German transliterations into corresponding phone sequences. This results in a dictionary that maps standard German words to Swiss pronunciations. Jointly with existing Swiss German lexicon resources (Schmidt et al., 2020), the previously generated mappings are then used to train a dedicated Swiss German G2P model.

We evaluate the quality of the resulting G2P model on a manually labeled test set. Those cover mappings from standard German words to Swiss German phone sequences and encompass a variety of relevant categories such as diminution, shortening or translation. Refer to Table 1 for samples of the assessed categories.

The Swiss G2P model allows to find suitable pronunciations for the relevant word inventory present in the acoustic and language model training corpora. However, to further increase the quality of the given pronunciations, data-driven lexicon learning techniques (Zhang et al., 2017) are applied. Those help to identify and correct noisy lexicon entries.

2nd person plural	2nd person sing	Diminuation	Shortening	Translation	Variability
fragt <i>fhr a_g ax t</i>	fragst <i>fhr a_k sh</i>	erdmännchen <i>e_r t m eh n l i_</i>	gymnasium <i>g ih m i_</i>	kopf <i>g hr ih n t</i>	kannst <i>k a sh</i>
riecht <i>sh m oe k c ax t</i>	riechst <i>sh m oe k sh</i>	gläschen <i>g l e_s l i_</i>	schwimmbad <i>b a_d ih</i>	kneipe <i>b ai ts</i>	zweites <i>ts v ai t</i>

Table 1: Example words and pronunciations from each G2P test condition

2.3 Language Model

Incorporating the translation from Swiss German to standard German as part of the lexicon introduces significant ambiguity in the decoding process. To counteract, we suggest using a strong standard German language model (LM) which helps to produce accurate hypotheses. We employ a first pass count-based LM to output up to 100 sentence hypotheses and a second pass neural LSTM (long short-term memory) LM (Sundermeyer et al., 2012) for rescoreing (Deoras et al., 2011). The first pass model is a 5-gram LM trained on large amounts of standard German text corpora totalling to over 100 billion words. We apply Kneser-Ney smoothing (Kneser and Ney, 1995).

Furthermore, we make some adjustments to better deal with Swiss German particularities, as described in the following paragraphs.

Compounds: German is a compounding language and tends to compose words (particularly nouns) of several smaller subwords. The resulting chains of word stems can lead to an infinitely large vocabulary size with words that occur very infrequently throughout the corpus. This spreading of probability mass weakens the LM. We hence decompound all compounded words in the training corpus and split them into subwords.

Clitics: Swiss German tends to merge words beyond compounding, not preserving word stems (Hollenstein and Aepli, 2014). For instance, the Swiss German ‘hemmer’ is the translation of ‘haben wir’ in standard German (English: ‘have we’). We identified approximately 8000 clitics in our corpus. We incorporate them in the decoding process by updating lexicon and LM. Following the example above the translated clitic ‘haben#wir’ with the corresponding Swiss pronunciation is added to the lexicon. As for the LM, we merge occurrences of relevant word pairs and interpolate with the unmerged LM.

2.4 Acoustic model

The acoustic model is trained with 80 dimensional log-mel filterbank features, with a processing window of 25ms and 10ms frame shift. The feature vector from the previous frame is concatenated with the current frame to obtain a 160 dimensional vector. We used a LC-BLSTM (latency controlled bidirectional long short-term memory) based acoustic model, that is popularly applied in speech recognition for controlling decoding latency to a few frames (Chen and Huo, 2016). The model was trained with alignments from a feed-forward network with context-dependent tied states (senones). The model has ~9k senone units. The LC-BLSTM is trained with 6 hidden layers with 512 units each. The hidden vectors from the forward and backward propagation were concatenated and then projected to a 512 dimensional vector. The model is trained with a cross entropy loss function. The decoding lexicon is extended with Swiss German words for training. The transliterations are used during forced alignment whenever possible. This helps to reduce the pronunciation ambiguity in the alignment phase and is especially helpful in the early training phases when no strong model is available for alignment.

The results are reported in terms of BLEU (Papineni et al., 2002) and word error rate (WER) on the Swiss Parliament test set described in Section 2.1.

3 Results and Discussions

An ablation study of the proposed approaches is presented in Table 2. All of the performance gains in this section will be reported as relative percentage improvements, while aforementioned table contains absolute numbers.

We first evaluate the effect of transfer learning on the results with the Swiss Parliament training set. It can be observed that it significantly helps to improve both WER and BLEU. In particular, the transfer-learned model (row 2, Table 2) improves over the model trained from scratch (row 1, Table 2) by 8.4% WER and 11.6% BLEU. The result shows

Row	Description	WER	BLEU
1	Swiss Parliament	45.93	33.16
2	+ Transfer Learning	42.09	37.53
3	+ Internal Data	41.10	38.49
4	+ 2 nd Pass Rescoring	38.71	41.91

Table 2: Performance in [%] of different system configurations evaluated on the Swiss Parliament test set.

that a well-trained German model can effectively boost the limited resources of Swiss German.

Further adding additional internal training data shows additional gains in performance. As such, we observe that the WER improves by 2.3% and BLEU by 2.5%.

Finally, 2nd Pass rescoring is applied as described in Section 2.3 to reorder the top 100 hypotheses. It can be observed from row 4, Table 2 that rescoring helps to improve the performance by 5.8% WER and 8.9% BLEU.

Our submission to SwissText 2021 achieves 46.04% BLEU on the official SwissText blind test set. This leads to a 12% relative margin in BLEU with respect to the second best competitor which was 40.99%.

The acoustic models have been trained using 8 GPUs for 25 epochs. This results in a total training time of around 400 GPU-hours when training on Swiss Parliament only and about 1200 GPU-hours when adding the internal data.

4 Conclusion and Future Work

In this paper, we described a speech recognition system that achieves strong results on the task of recognizing Swiss German dialect and translating it into standard German text. We proposed a hybrid ASR system with a lexicon that incorporates translations, a 1st pass language model that deals with Swiss German word compounding and clitics, an acoustic model that is transfer-learned from standard German resources and a strong neural language model for 2nd pass rescoring to smoothen translation artifacts. Furthermore, we provided an ablation study that allows to infer the effect of adding training data, performing transfer learning and 2nd pass rescoring. Our submission reached 46.04% BLEU on a challenging conversational test set and outperformed all competing approaches by a large margin.

In terms of future work, we would like to investigate word re-orderings as part of the translation, which our current model does not actively

support. For instance, Swiss German frequently moves verbs in relative clauses to different positions with respect to the standard German word order. Furthermore, sequence discriminative training is a promising route for exploration as well as using unsupervised data for acoustic model training.

References

- Aashish Agarwal and Torsten Zesch. 2020. Ltl-ude at low-resource speech-to-text shared task: Investigating mozilla deepspeech in a low-resource setting.
- Ahmed Ali, Shammur Chowdhury, Mohamed Afify, Wassim El-Hajj, Hazem Hajj, Mourad Abbas, Amir Hussein, Nada Ghneim, Mohammad Abushariah, and Assal Alqudah. 2021. [Connecting Arabs: Bridging the gap in dialectal speech recognition](#). *Communications of the ACM*, 64(4):124–129.
- Ahmed Mohamed Abdel Maksoud Ali. 2018. *Multi-Dialect Arabic Broadcast Speech Recognition*. Ph.D. thesis, University of Edinburgh, Edinburgh, UK.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Hervé Bourlard and Stéphane Dupont. 1996. A new ASR approach based on independent processing and recombination of partial frequency bands. In *Proc. Int. Conf. on Spoken Language Processing (ICSLP)*, volume 1, pages 426–429.
- 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.
- Kai Chen and Qiang Huo. 2016. [Training deep bidirectional lstm acoustic model for lvcsr by a context-sensitive-chunk bptt approach](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24(7):1185–1193.
- Anoop Deoras, Tomáš Mikolov, and Kenneth Church. 2011. A Fast Re-scoring Strategy to Capture Long-Distance Dependencies. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1116–1127, Edinburgh, Scotland, UK. Association for Computational Linguistics.

- Philip N. Garner, David Imseng, and Thomas Meyer. 2014. Automatic Speech Recognition and Translation of a Swiss German Dialect: Walliserdeutsch. In *Proc. of the Ann. Conf. of the Int. Speech Commun. Assoc. (Interspeech)*, Singapore.
- Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, et al. 2014. Deep speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567*.
- Nora Hollenstein and Noëmi Aeppli. 2014. Compilation of a Swiss German dialect corpus and its application to PoS tagging. In *Proceedings of the First Workshop on Applying NLP Tools to Similar Languages, Varieties and Dialects*, pages 85–94.
- Amir Hussein, Shinji Watanabe, and Ahmed Ali. 2021. [Arabic Speech Recognition by End-to-End, Modular Systems and Human](#). *arXiv:2101.08454 [cs, eess]*.
- Tannon Kew, Iuliia Nigmatulina, Lorenz Nagele, and Tanja Samardzic. 2020. Uzh tilt: A kaldi recipe for swiss german speech to standard german text.
- Reinhard Kneser and Hermann Ney. 1995. Improved backing-off for m-gram language modeling. In *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, volume 1, pages 181–184.
- Jason Li, Vitaly Lavrukhin, Boris Ginsburg, Ryan Leary, Oleksii Kuchaiev, Jonathan M Cohen, Huyen Nguyen, and Ravi Teja Gadde. 2019. Jasper: An end-to-end convolutional neural acoustic model. *arXiv preprint arXiv:1904.03288*.
- Iuliia Nigmatulina, Tannon Kew, and Tanja Samardzic. 2020. [ASR for non-standardised languages with dialectal variation: the case of Swiss German](#). In *Proceedings of the 7th Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 15–24, Barcelona, Spain (Online). International Committee on Computational Linguistics (ICCL).
- 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. Association for Computational Linguistics.
- Michel Plüss, Lukas Neukom, and Manfred Vogel. 2020. Swiss parliaments corpus, an automatically aligned swiss german speech to standard german text corpus. *arXiv preprint arXiv:2010.02810*.
- Larissa Schmidt, Lucy Linder, Sandra Djambazovska, Alexandros Lazaridis, Tanja Samardžić, and Claudiu Musat. 2020. [A Swiss German Dictionary: Variation in Speech and Writing](#). *arXiv:2004.00139 [cs]*.
- Michael Stadtschnitzer and Christoph Schmidt. 2018. Adaptation and training of a swiss german speech recognition system using data-driven pronunciation modelling. *Proceedings of DAGA-44. Jahrestagung für Akustik, München, Germany*.
- Martin Sundermeyer, Ralf Schlüter, and Hermann Ney. 2012. LSTM neural networks for language modeling. In *Proc. of the Ann. Conf. of the Int. Speech Commun. Assoc. (Interspeech)*, pages 194–197, Portland, OR, USA.
- Wayne Xiong, Jasha Droppo, Xuedong Huang, Frank Seide, Mike Seltzer, Andreas Stolcke, Dong Yu, and Geoffrey Zweig. 2016. Achieving human parity in conversational speech recognition. *arXiv preprint arXiv:1610.05256*.
- Xiaohui Zhang, Vimal Manohar, Daniel Povey, and Sanjeev Khudanpur. 2017. [Acoustic Data-Driven Lexicon Learning Based on a Greedy Pronunciation Selection Framework](#). In *Proc. of the Ann. Conf. of the Int. Speech Commun. Assoc. (Interspeech)*, pages 2541–2545. ISCA.