

Multi-Staged Cross-Lingual Acoustic Model Adaption for Robust Speech Recognition in Real-World Applications — A Case Study on German Oral History Interviews

Michael Gref,^{1,2} Oliver Walter,¹ Christoph Schmidt,¹ Sven Behnke,^{1,3} Joachim Köhler¹

¹Fraunhofer Institute for Intelligent Analysis and Information Systems (IAIS), Germany

²Institute for Pattern Recognition (iPattern), Niederrhein Univ. of Applied Sciences, Germany

³Autonomous Intelligent Systems (AIS), Computer Science Institute VI, Univ. of Bonn, Germany

{michael.gref, oliver.walter, christoph.andreas.schmidt, sven.behnke, joachim.koehler}@iais.fraunhofer.de

Abstract

While recent automatic speech recognition systems achieve remarkable performance when large amounts of adequate, high quality annotated speech data is used for training, the same systems often only achieve an unsatisfactory result for tasks in domains that greatly deviate from the conditions represented by the training data. For many real-world applications, there is a lack of sufficient data that can be directly used for training robust speech recognition systems. To address this issue, we propose and investigate an approach that performs a robust acoustic model adaption to a target domain in a cross-lingual, multi-staged manner. Our approach enables the exploitation of large-scale training data from other domains in both the same and other languages. We evaluate our approach using the challenging task of German oral history interviews, where we achieve a relative reduction of the word error rate by more than 30% compared to a model trained from scratch only on the target domain, and 6–7% relative compared to a model trained robustly on 1000 hours of same-language out-of-domain training data.

Keywords: acoustic modeling, acoustic model adaption, cross-lingual, digital humanities, oral history, speech recognition, transfer-learning, under-resourced speech recognition

1. Introduction

Current automatic speech recognition (ASR) systems show remarkable performance on tasks where large amount of annotated and time-aligned speech data is used for training. Other tasks, however, where only little or no data is present—such as under-resourced or zero resource languages—are still a challenge for modern systems. Nevertheless, also within one language, the recognition performance drastically degrades if off-the-shelf systems are applied to different domains in terms of recording conditions and speech characteristics, such as noisy and reverberated recordings, spontaneous, fast speech, accents, unclear pronunciations and age- and health related changes in the ways of speaking—to name but a few. In this work, we propose and investigate an approach that performs a robust acoustic model adaption to a target domain in a cross-lingual, multi-staged manner. The proposed approach utilizes large-scale training data from rich-resourced domains in both the same and other languages.

A lot of investigations in the field of speech recognition are based on well-known, publicly available benchmark data sets like Switchboard, Wall Street Journal and AMI to evaluate proposed systems and approaches. Usually, such data sets cover well-defined, restricted challenges and are of good use to study this specific challenge. However, they are only of limited use for replicating the conditions of challenging real-world applications where many of the aforementioned issues might be present at the same time with varying degrees.

German *oral history* interviews are such a challenging real-world application for automatic speech recognition in which the aforementioned issues are present to an unknown extent and with unknown distribution in each sample. The term oral history, in historical research, refers to conducting

and analyzing interviews with contemporary witnesses.

Particularly in Germany, this kind of research focuses above all on the period of the Second World War and National Socialism. But in the meantime, it has also come to include many other topics and historical periods. The past forty years have seen a multitude of witnesses to a wide range of historical events interviewed by researchers. In the original sense oral history interviews are characterized by the fact that rather than structuring the interview around questions, the interviewer encourages the interviewee to freely narrate his or her life story. In terms of biographical research, the outcome is qualified as a narrative life-story interview lasting very often at least three or four hours.

For both, archiving and analyzing, time-aligned transcription and indexing with keywords of such interviews is essential. Due to the poor performance of off-shelf recognition systems on such interviews and the lack of sufficient in-domain training data, transcribing, labeling and annotating speech recordings is still often performed manually. The huge effort in terms of time and human resources required to do this, severely limits the efficiency to exploit oral history interviews for digital humanities research.

As our contribution for solving this problem we developed and studied an approach on this unconstrained, demanding use-case. Further, we investigate the performance of models trained with the proposed approach using German test sets from several other domains. In this way we demonstrate the generalization and robustness of the model for similar related challenges in other domains and the applicability of the system in real applications. In addition, our ablation study experiments show that the cross-lingual adaption of state-of-the-art LF-MMI acoustic models results in significantly improved speech recognition performance—

even if only a very small amount of training data from the target domain is used. This observation can contribute to the ongoing research on speech recognition for under-resourced language.

2. Related Work

The use of automatic speech recognition technology to transcribe and index oral history interviews has started with the MALACH project (Pstuka et al., 2002) where interviews of the *Survivors of the Shoah Visual History Foundation* (VHF) were processed with a state-of-the-art speech recognition in 2002. Oard (2012) investigated how speech recognition technology can be used for oral history research. In our prior work (Gref et al., 2018a) we studied the application and adaption of state-of-the-art speech recognition systems to German oral history interviews. Furthermore in that work we proposed an ASR test set for German oral history interviews that we utilize in this work for evaluation. Just recently, Picheny et al. (2019) proposed parts of data from the MALACH project as a speech recognition challenge for English oral history. The authors agree with our statements, that the challenges of such interviews are still open problems for modern speech recognition systems. In contrast to the German oral history challenge, the MALACH corpus provides a few hundred hours of annotated in-domain training data. For German oral history, such amounts of data are currently not available. Furthermore, in contrast to the German oral history evaluation set used in our work, the interview in the MALACH project are recorded in relatively high-quality audio quality. The German oral history interviews in our evaluation set, however, were conducted by historians throughout many decades with a wide range of recording equipment, setups, and acoustic conditions.

ASR is a popular and highly researched area and new approaches are regularly proposed. End-to-end approaches to speech recognition have received a lot of attention lately. Since these models do not require alignments from a bootstrap model—as conventional HMM-DNN approaches do—these models are significantly simpler to train. In many scenarios, however, conventional HMM-DNN systems perform better than end-to-end systems. This is especially the case when only a small amount of training data is used. In ASR, lattice-free maximum mutual information (LF-MMI) trained models (Povey et al., 2016) achieved state-of-the-art results on many different speech recognition tasks in recent years and are applied in many recently proposed speech recognition systems.

A common approach for certain out-of-domain tasks is to apply data augmentation on rather clean data and perform multi-condition training. Ko et al. (2017) studied data augmentation of reverberant speech for state-of-the-art LF-MMI models. Speed perturbation techniques to increase training data variance have been investigated by Ko et al. (2015). The proposed method in this work is to increase the data three-fold by creating two additional versions of each signal using the constant speed factors 0.9 and 1.1—a method used in many recent training routines by default. In (Gref et al., 2018b) we studied noise and reverberation data augmentation for robust speech recognition on German oral history interviews for conventionally (cross-

entropy) trained and LF-MMI models.

While data augmentation is quite successful to overcome a mismatch in acoustic conditions between desired applications and training data, it is limited to acoustic distortions that can be artificially created. In recent years, transfer learning for acoustic model adaption has raised attention for under-resourced language tasks. Transfer learning is an approach to improve generalization and performance by transferring knowledge of a model trained in one scenario to train a model in another related scenario (Goodfellow et al., 2016). It is particularly useful in scenarios where only little training data is available for the main task but a large amount of data is available for a similar or related task.

Wang and Zheng (2015) give a detailed overview of transfer learning in speech and language processing. Ghahremani et al. (2017) investigated transfer learning using LF-MMI models for many different well-known English speech recognition tasks. For the German oral history task, we studied transfer learning using a very small amount of training data and a leave-one-speaker-out evaluation approach (Gref et al., 2019).

Cross-lingual acoustic model adaption or knowledge transfer aims at utilizing the knowledge of models trained in one or more languages to improve speech recognition performance for a target language. Xu et al. (2016) studied semi-supervised learning and cross-lingual knowledge transfer with multilingual data and neural network fine-tuning. Feng and Lee (2018) investigated cross-lingual knowledge transfer in a multilingual setup with language-dependent pre-final layers under each softmax output layer. Ma et al. (2017) have studied multilingual training using LF-MMI models with a joint output layer across languages followed by the adaption to a low-resourced target language.

3. Proposed Approach

The proposed approach consists of three proceeding stages of acoustic model adaption, as shown in Figure 1. Throughout the stages, we first use a large amount of data from more far-away domains and then smaller amounts of data from nearby domains for training. At the transition of each stage, we transfer the corresponding learned knowledge to the network in the next stage.

3.1. Data Augmentation

To achieve robustness to acoustic conditions of the models trained in the first two stages, we apply noise and reverberation data augmentation increasing the data threefold. Defining discrete-time-signals as sequences of sample values, the applied augmentation can be described as

$$(x_n)_{n \in \mathbb{N}} := (s_n)_{n \in \mathbb{N}} * (h_n)_{n \in \mathbb{N}} + (w_n)_{n \in \mathbb{N}} * (\tilde{h}_n)_{n \in \mathbb{N}} \quad (1)$$

if both noise and reverberation inside a simulated room affects the speech signal. Here, $*$ is the convolution operation for sequences, $(s_n)_{n \in \mathbb{N}}$ the sequence of the clean speech signal, $(h_n)_{n \in \mathbb{N}}$, $(\tilde{h}_n)_{n \in \mathbb{N}}$ are room impulse responses modeling the reverberation of one room at different positions and $(w_n)_{n \in \mathbb{N}}$ describes the sequence of the noise signal.




Figure 1: Proposed approach for cross-lingual, multi-staged acoustic model adaption.

If only reverberation and no background noise affects the speech signal, $\forall n \in \mathbb{N} : w_n = 0$ applies and yields

$$(x_n)_{n \in \mathbb{N}} := (s_n)_{n \in \mathbb{N}} * (h_n)_{n \in \mathbb{N}}. \quad (2)$$

3.2. Stage 1: Training on Huge Amount of English Data

In the first stage of the proposed approach, as shown in Figure 1, a robust acoustic model is trained using a huge amount of different-language, out-of-domain training data. A model that is well trained on a large amount of data has learned to perform a robust extraction of relevant acoustic input features and learned useful internal representations for the classification task. We assume that these aspects are, at least to some extent, language independent for related languages. To use this knowledge for tasks in languages with less available data, we apply a weight transfer of all hidden layers and use these layers to initialize the acoustic model neural network training in the second stage. Currently, probably English is the language with the largest amount of available training data for speech recognition.

Therefore, we propose combining several English corpora from different domains to train the acoustic model in this stage. Furthermore, in our approach we apply the aforementioned noise and reverberation data augmentation to further increase the robustness and generalization of the model increasing training data three-fold.

As a default step of the acoustic model training, we train an i-vector extractor on the English data in this stage. We further use this i-vector extractor in the following two stages for the acoustic model adaption training. Currently, we do not retrain the extractor on data from the target language.

3.3. Stage 2: Transfer Learning on German (Out-Of-Domain) Broadcast Data

In the second stage, the model is adapted to the language of the target domain, however, with training data from another domain. We perform the cross-lingual adaption implicitly by replacing the language-dependent output layer of the neural network trained on English with a randomly initialized output layer for German tied states—and then train

this model on German data. No manual phoneme mapping is applied.

In this stage, the feature extraction and representation learning of the lower layers is further improved for the target language. The classification of subphonetic units, which are language dependent, is learned in the output layer. For our task, we use 1000 hours of German broadcast data in this stage. Again, in this stage training data is increased three-fold using noise and reverberation data augmentation with slightly different configuration than in the previous stage.

The training setup in this stage is completely equal to the one used in the previous stage. Parameter tuning, such as adapting the learning rate layer-wise, might lead to better results. However, due to the large computational load to train a state-of-the-art acoustic model neural network on 1000 hours, we leave this for future work.

3.4. Stage 3: Transfer Learning on Automatically Aligned German Oral History Interviews

In the last stage, the acoustic model is finally adapted to the target domain. We do not only transfer the hidden layers in this stage but we apply a full weight transfer of the entire source model for initialization of the target model. In particular, we do not replace the output layer in this stage, since we use the same set of phonemes and the same phonetic decision tree both in Stage 2 and Stage 3. We obtain the respective training lexicons using the same grapheme-to-phoneme (G2P) pronunciation model.

We apply no dropout in this transfer learning stage. In our experiments, we have not found a dropout setup that has led to a notable improvement using very little training data. The initial learning rate in this stage is 100 times smaller than the final learning rate used for the training in the previous stages. These values are our default learning rates based on our previous experiments. All other training parameters are kept equal in all stages.

Since, we have lack of training data for German oral history interviews, we apply automatic audio transcript alignment according to Manohar et al. (2017), but with neural network based acoustic models trained on German broadcast data, to semi-automatically create in-domain training from several transcribed but not time-aligned oral history recordings.

We consider the transcriptions used to create the data in this stage partly erroneous and incomplete, since they are automatically parsed from several source data to normalized plain text. In our task, these are interview transcriptions by different historians using different styles for annotation, containing remarks and notes in the middle of the text and various document types, such as MS Word and RTF.

We perform automatic text normalization to write out abbreviations, numbers and dates, and restore casing of words at the beginning of sentences as good as possible. However, we generally perform no manual correction of transcripts or selection of data for alignment. The acoustic model we used to create these alignments is a slightly weaker LF-MMI model from prior works, trained robustly on the 1000 hours of German broadcast data that are also used in Stage 2.

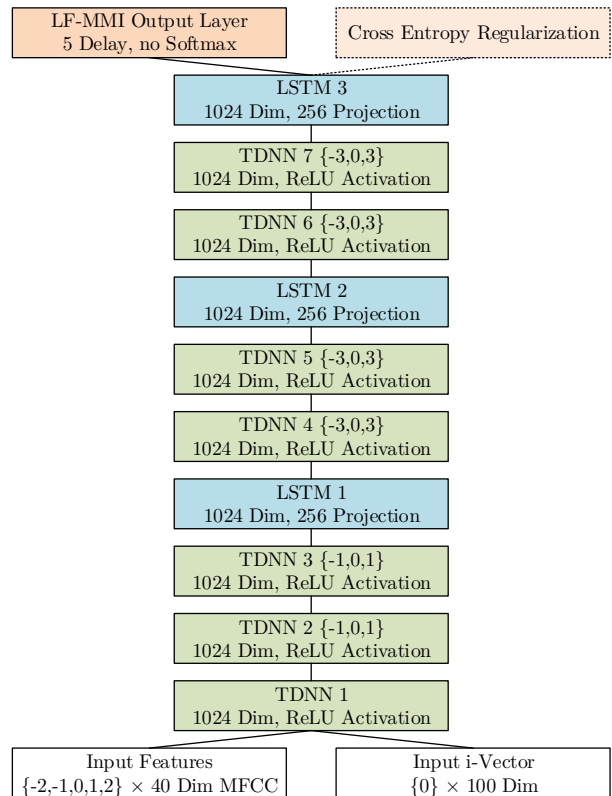


Figure 2: Architecture of the acoustic model neural network used for our experiments.

4. Experimental Setup

We carry out all experiments using the Kaldi ASR toolkit (Povey et al., 2011). Neural network acoustic models are trained using the aforementioned LF-MMI (Povey et al., 2016) approach and speed perturbation (Ko et al., 2015) increasing the amount of data three-fold using constant speed factors 0.9 and 1.1.

4.1. Acoustic Model Neural Network

The acoustic model in our experiments uses a 300-dimensional input at each time-step consisting of five consecutive 40-dimensional MFCC features and a 100-dimensional i-vector (Dehak et al., 2011). We use a topology with ten hidden layers that was proposed and investigated by Cheng et al. (2017) with seven TDNN layers (Waibel et al., 1989), (Peddinti et al., 2015) and three LSTM layers (Hochreiter and Schmidhuber, 1997) stacked in the order given in Figure 2. The applied implementation uses LSTM layers with forget gates (Gers et al., 2000), peephole connections (Gers and Schmidhuber, 2000) and projection layers (Sak et al., 2014). The LSTM layers have a cell dimension of 1024 and a projection dimension of 256. The TDNN layers are 1024-dimensional as well.

4.2. Training of English Model in Stage 1

For training the English model in Stage 1, we combined the training data sets from the well-known corpora *Librispeech* (Panayotov et al., 2015), *Commonvoice*¹, *Switchboard* (Godfrey et al., 1992) and *Fisher* (Cieri et al., 2004).

¹voice.mozilla.org

Table 1: Statistics of the used evaluation sets

Evaluation Set	Length [min.]	∅Segment Length	∅Words Per Segment	∅Words Per Second
Oral History	211	5.3 s	11.6	2.2
Interaction	49	1.1 s	3.8	3.4
Spoken QALD-7	15	4.3 s	6.9	1.6
Challenging Broadcast	105	10.6 s	29.3	2.8
DiSCo (Broadcast Eval)	323	2.4 s	6.9	2.9

Overall, the English data used in training comprises more than 3400 hours of annotated speech. Additionally to the clean data set, we create two distorted versions of the data using noise and reverberation data augmentation according to Equation 1. The first version uses a signal-to-noise ratio between 5 and 10 dB, the second one between 10 to 20 dB. In both versions, we utilize 266 room impulse responses of small and medium-sized rooms for reverberation and several noises recorded in real-life scenarios.

Using a general propose language model for decoding that is based on crawled data, the English model achieves 9.17% word-error-rate on Librispeech test-clean and 18.30% on the data from *Voices Obscured in Complex Environmental Settings (VOICES)* (Richey et al., 2018).

4.3. Adaption to German in Stage 2

To adapted the English model from Stage 1 to German, we utilize the 1000-hour large-scale corpus of German broadcast speech *GerTV1000h* (Stadtshnitzer et al., 2014) in this stage. We considered this data out of domain for the oral history task, since broadcast recordings differ from oral history in terms of the used recording technology, audio signal quality and speech characteristics. Broadcast recordings are generally recorded in clean conditions using professional equipment and expertise. In most cases, the speech is clearly pronounced, well articulated and easy to understand. In contrast, oral history interviews are often recorded using conventional recording equipment that was common at the time of recording. The speakers often have a large distance to the microphone and utter sentences spontaneously as they come to their mind. A characteristic attribute of many oral history interviews are the age- and health-related changes in the way elderly people speak.

To perform the noise and reverberation data augmentation in this stage, two additional distorted versions of the clean training are created. In the first version, we apply artificial reverberation according to Equation 2 using 266 room impulse responses of small and medium-sized rooms. The second one is created according to Equation 1 with both reverberation and noises using a random signal-to-noise ratio between 10 and 20 dB.

4.4. Adaption to Oral History Target Domain in Stage 3

Randomly selected oral history interviews of 150 contemporary witnesses serve as data from the target domain for the automatic transcript alignment and domain adaption in Stage 3. It is ensured that only interview of speakers that do not appear in the test data are used for training.

The recording quality differs from barely understandable to good quality with only light distortions. These are mostly low energy noises from cassette recording machines or reverberation due to a large distance to the recording microphone in a medium-sized room. Some speakers have a slight dialect or accent, but most speak High German.

Different historians throughout many years transcribed these interviews using varying annotation styles and formats. Since we performed no manual check, we could not guarantee that the entire transcription of the interview is present for alignment. For the final training we use 25 hours of automatically aligned speech.

4.5. Evaluation Sets

We do not only report results on the aforementioned oral history test set but overall on four additional German in-house test sets from different domains. Some of these sets partly share some challenges of oral history interviews, such as spontaneous speech with unclear pronunciations, noises and reverberation. The sets used in our work are listed in Table 1 along with some statistics characterizing the sets. Along with the overall test set length, the average segment length, the average amount of words per segment, and the average spoken words per second are reported. Very short word sequences per segment usually tend to benefit less from a large powerful language than segments with rather long word sequences. A high number of spoken words per second in a test set indicates a fast speaking speed. The faster speakers are speaking, the more difficult it gets for the speech recognition system—as it gets for human listeners as well—to distinguish different phonemes and words.

As evaluation set for the target domain, we use the German oral history data set from (Gref et al., 2018a). It consists of 3.5 hour audio from 35 different speakers recorded in real oral history interviews. The recordings took place between 1980 and 2012, representing a wide range of recording technology, interview methodology, dialects and pronunciations. The set is manually transcribed and segmented.

The *Interaction* evaluation set contains recordings of people informally talking to each other recorded in challenging acoustic conditions. This set is characterized by very fast, partly overlapping, highly colloquial speech, speaker noises such as laughter and unclear pronunciation. This set has the highest average speed of speaking of all used test sets. Furthermore, due to the nature of fast, informal conversations, the average segment length is the shortest and many segments only contain one or two words. This makes this test set one of the most challenging evaluations sets in

Table 2: Word error rates of the proposed approach compared to two baselines and the ablation study. Results are reported both for the default and the large decoding language model.

		Baseline Broadcast	Baseline Oral History	Removing Stage 1	Removing Stage 2	Removing Stage 3	Proposed Approach
Setup	Stage 1 (English)				×	×	×
	Stage 2 (German Broadcast)	×		×		×	×
	Stage 3 (Ger. Oral History)		×	×	×		×
Evaluation Set	Oral History	27.66	37.42	25.93	28.67	27.38	25.91
	+ larger decode LM	27.08	38.18	25.26	28.48	26.48	25.17
	Interaction	48.23	69.07	48.17	58.60	47.35	47.14
	+ larger decode LM	51.22	71.96	51.36	60.13	50.58	50.26
	Spoken QALD-7	18.95	36.74	19.09	31.08	18.61	18.40
	+ larger decode LM	14.79	30.06	14.25	24.61	13.77	13.57
	Challenging Broadcast	19.68	33.55	19.81	26.58	19.47	19.38
	+ larger decode LM	17.42	31.90	17.61	24.59	17.34	17.36
DiSCo (Broadcast Eval)	11.89	27.91	12.45	20.81	11.85	12.37	
+ larger decode LM	12.15	28.73	12.48	20.75	12.12	12.42	

our data collection.

The *Spoken QALD-7* corpus contains in-house recorded questions for a question answering system described in (Usbeck et al., 2017), where several speakers used a web interface and their respective microphone—a headset or build-in laptop microphone for instance. One segment usually comprises exactly one question prompt.

The *Challenging Broadcast* evaluation set contains a 1.75 hour collection of rather challenging interviews and recordings from the German broadcast domain with a lot of spontaneous speech, often in challenging acoustic conditions, including overlapping and dialectal speech. The average speaking speed is somewhat between the oral history and the interaction test set.

DiSCo (Baum et al., 2010) is a corpus for the German broadcast domain and contains four evaluation sets: *planned* and *spontaneous* speech, each in *clean* and *mixed* acoustic conditions. For the sake of clarity, we always report the unweighted mean value of these subsets.

4.6. Lexicon and Language Model

The training and decoding lexicons are obtained using a G2P model trained with Sequitur G2P (Bisani and Ney, 2008). This model is created using the German pronunciation dictionary *Phonolex* from the *Bavarian Archive for Speech Signals*.

To assure the improvements of the acoustic model adaptation using our proposed approach are consistent and do not depend on a certain language model, we perform two independent decodings in each experiment using two different language models. First, we use our default 500,000 words broadcast language model. This model is trained on broadcast text corpora consisting of 75 million words. Further, we use a larger broadcast language model that is trained on texts with 1.6 billion running words of crawled German news data websites. The lexicon of this model contains more than 2 million different words. Both models are trained as 5-gram language models with a 1e-8 pruning factor. All decoding parameters are kept to fix default values for all of our experiments. Especially, the language model

weight is kept to a fixed value for all experiments.

4.7. Performed Experiments

We use two baseline models to compare our proposed approach. These models are trained without neural network initialization from a prior model and not further adaption is performed. As the first baseline serves a model trained from scratch on the 1000 hours of German broadcast data with the same configuration and setup as in Stage 2. The second baseline is a model that we train from scratch on the 25 hours of automatically aligned oral history data. Furthermore, to determine which of the stages contributes to the improvements to what extent, we carry out ablation study experiments in which we omit one of the stages from training.

In each experiment in which the English training in Stage 1 is skipped (both in the baselines and in the ablation study), the i-vector extractor is also trained with the corresponding data used in the first respective training step. To better understand, where the increased robustness through knowledge transfer from Stage 1 comes from, we also perform some experiments comparing the German and English trained i-vector extractor.

5. Results

The results of the proposed approach, baseline models, and the ablation study are summarized in Table 2. For each test set, the upper row shows the achieved word error rate using the default 500,000 words language model. The respective next row shows the results with the larger 2-million-words language model.

5.1. Comparison to Baselines

The oral history baseline performs significantly worse than the broadcast baseline on all sets—even on the target domain. This is due to the significantly smaller amount of training data. Compared to the broadcast baseline, the proposed approach achieves a relative word error improvement of 6.3% on the oral history test set using the default lan-

guage model. Using the larger language model, the relative improvement is 7.1%.

On all other evaluation sets, except DiSCo, we observe an improvement using the proposed approach. This indicates a good generalization of the model and suggests that the model is robust enough for operation in real-world applications. The slightly reduced performance on the DiSCo set is probably because this test data is already very close to the conditions presented by the broadcast training used to train the baseline model.

5.2. Ablation Study

5.2.1. Removing Stage 1

In the first ablation study setup, we remove the adaption from the English model trained in Stage 1. In particular, we randomly initialize the model in Stage 2 for training on German broadcast data and then adapted to the oral history data in Stage 3. Furthermore, we train a new i-vector extractor on German data. Compared to the broadcast baseline, we observe a significant word error rate improvement using both language models on the oral history test set. The achieved word error rate on this test set is almost as good as offered by the proposed approach. However, compared to the broadcast baseline, on the Challenging Broadcast and DiSCo test set, the results get worse with both language models in this setup. This is also true for the QALD-7 test set decoded with the smaller language model and for the interaction test set and the larger language model. This indicates that the initialization with the English trained model increases the robustness of the proposed approach for many domains.

5.2.2. Removing Stage 2

In this setup, we skip the adaption to German broadcast in Stage 2 and instead directly adapt the English trained model both to the target language and target domain using the oral history training data. While the results are still worse than the broadcast baseline, a remarkable improvement is achieved compared to the oral history baseline. Using only English training data and 25 hours of automatically aligned German training data, we achieve a word error rate reduction from 37.4 to 28.7% on the target domain decoding with the default language model. Results are similar for the large language model.

This value is only one absolute percentage point higher than the value achieved by the broadcast baseline that we trained on 1000 hours of manually annotated German data. Thus, adapting from a rich-resourced language directly to the target language and domain is a reasonable approach if no other data is available for training in the target language. In particular, this is shown in this experiment for state-of-the-art LF-MMI acoustic models with an LSTM-TDNN topology, which usually require a lot of data during training. For real-world applications one cannot expect the model to be as robust as models trained robustly on large-scale data from the target language—but definitely more robust than training from scratch on the small data set only which can be usefully especially for under-resourced languages.

Table 3: Comparison of German and English trained i-vector extractors in acoustic model training on German broadcast data with random model initialization and with knowledge transfer from English-trained model.

Setup	Ac. Model Init. i-vector	Rand. Ger.	Eng. Ger.	Rand. Eng.	Eng. Eng.
Eval Set	Oral History	27.66	27.75	27.70	27.38
	Interaction	48.23	48.54	47.83	47.35
	Spok. QALD-7	18.95	19.15	19.43	18.61
	Chall. Broadcast.	19.68	19.78	19.57	19.47
	DiSCo	11.89	11.86	11.88	11.85

5.2.3. Removing Stage 3

In the last setup of the ablation study, we omit the fine-tuning to target domain and evaluate the model trained on broadcast data in Stage 2. Compared to the broadcast baseline, this experiment shows once again that initialization with the English-trained model leads to an increase in robustness and better results on all test sets. Remarkably, this model even achieves the best results on the DiSCo test set which we consider close to the German broadcast training data. This setup also achieves slightly better results than the proposed approach on the Challenging Broadcast test set with the larger language model. However, the gain is less than 0.2% relative.

5.3. Influence of the i-Vector Extractor

In the following experiments, we study whether the improved robustness through knowledge transfer of the English model is due to weight transfer, as we expect, or due to better i-vector extractor trained on English data instead of German. Therefore, we compare each combination of using German or English trained i-vector extractors each with a random model initialization or initialization of hidden layers with the English trained model when training on the 1000 hour German broadcast data. For simplicity, we only report results on the small language model. The results are summarized in Table 3.

Using the English-trained i-vector extractor along with acoustic model initialization from the English-trained model—as is the case in our proposed approach—we achieve the best results on all test sets. Using the German i-vector extractor with English model initialization leads to the worst result for three out of the five test sets. This is to be expected since i-vectors of similar speakers from two differently trained i-vector extractors point to completely different directions in the two different 100-dimensional vector spaces. Since the English model is trained with i-vectors from one space, using different i-vectors in the second stage causes wrong estimations of speakers and this relation has to be relearned for the new vector space in Stage 2.

Using the English i-vector extractor instead of the German one with random acoustic model initialization only leads to slightly better results on three test sets and slightly worse results on two test sets. Therefore, we conclude that, in deed, weight transfer leads to better results—and not a bet-

ter i-vector extractor trained on English data. However, in order to use weight transfer sensibly, the i-vector extractor must be used, which was also used to train the original model.

6. Conclusions

In this work, we proposed and investigated an approach that performs a robust acoustic model adaption to a target domain in a cross-lingual, multi-staged manner. Our approach addresses challenges where only little training data for the target domain are present and enables the exploitation of large-scale training data from other domains in both the same and other languages. These other domains may be acoustic distortions such as reverberation or poor recording equipment, but also deviations in the way of speaking, articulation, dialects, spontaneous speech, unclear articulations and much more.

We studied our approach using the challenging example of German oral history interviews with the aim to obtain an acoustic model that is robust enough to be applied in real-world applications. In our experiments we first trained a robust acoustic model for English with more than 3000 hours of data, that we then adapted to German using 1000 hours of German broadcast data. This model is again adapted using 25 hours of in-domain oral history interview.

To determine the robustness of the model, for this case study we performed exhausting experiments and evaluated not only on in-domain data but also on several German test sets from other domains and two different decoding language models. To thoroughly determine which stage of the proposed approach leads to which improvement, we further performed different ablation study experiments. Thereby we have shown that the direct adaptation of LF-MMI acoustic models from one language to another leads to good results, even using only very little training data from the target domain. Thus, this observation can contribute to the ongoing research on speech recognition for under-resourced language.

The model trained with our proposed approach achieves a relative reduction of the word error rate by more than 30% compared to a model trained from scratch only on the target domain, and 6–7% relative compared to a model trained robustly on 1000 hours German broadcast training data.

7. Acknowledgements

This research has been funded by the Federal Ministry of Education and Research of Germany (BMBF) in the project KA³ - *Kölner Zentrum für Analyse und Archivierung von AV-Daten* (Cologne center for the analysis and archiving of audiovisual data) (project number: 01UG1811B).

8. Bibliographical References

Bisani, M. and Ney, H. (2008). Joint-sequence models for grapheme-to-phoneme conversion. *Speech Communication*, 50(5):434–451.

Cheng, G., Peddinti, V., Povey, D., Manohar, V., Khudanpur, S., and Yan, Y. (2017). An exploration of dropout with lstms. In *18th Annual Conference of the International Speech Communication Association (Interspeech)*, pages 1586–1590.

Dehak, N., Kenny, P., Dehak, R., Dumouchel, P., and Ouellet, P. (2011). Front-end factor analysis for speaker verification. *IEEE Trans. Audio, Speech & Language Processing*, 19(4):788–798.

Feng, S. and Lee, T. (2018). Improving cross-lingual knowledge transferability using multilingual TDNN-BLSTM with language-dependent pre-final layer. In *19th Annual Conference of the International Speech Communication Association (Interspeech)*, pages 2439–2443.

Gers, F. A. and Schmidhuber, J. (2000). Recurrent nets that time and count. In *IJCNN (3)*, pages 189–194.

Gers, F. A., Schmidhuber, J., and Cummins, F. A. (2000). Learning to forget: Continual prediction with LSTM. *Neural Computation*, 12(10):2451–2471.

Ghahremani, P., Manohar, V., Hadian, H., Povey, D., and Khudanpur, S. (2017). Investigation of transfer learning for ASR using LF-MMI trained neural networks. In *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 279–286.

Goodfellow, I. J., Bengio, Y., and Courville, A. C. (2016). *Deep Learning*. Adaptive computation and machine learning. MIT Press. pp. 526–528.

Gref, M., Köhler, J., and Leh, A. (2018a). Improved transcription and indexing of oral history interviews for digital humanities research. In *Eleventh International Conference on Language Resources and Evaluation (LREC)*, pages 3124–3131.

Gref, M., Schmidt, C., and Köhler, J. (2018b). Improving robust speech recognition for german oral history interviews using multi-condition training. In *13. ITG Symposium on Speech Communication*, pages 256–260. IEEE.

Gref, M., Schmidt, C., Behnke, S., and Köhler, J. (2019). Two-staged acoustic modeling adaption for robust speech recognition by the example of german oral history interviews. In *IEEE International Conference on Multimedia and Expo, ICME, Shanghai, China*, pages 796–801.

Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8):1735–1780.

Ko, T., Peddinti, V., Povey, D., and Khudanpur, S. (2015). Audio augmentation for speech recognition. In *16th Annual Conference of the International Speech Communication Association (Interspeech)*, pages 3586–3589.

Ko, T., Peddinti, V., Povey, D., Seltzer, M. L., and Khudanpur, S. (2017). A study on data augmentation of reverberant speech for robust speech recognition. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5220–5224.

Ma, J. Z., Keith, F., Ng, T., Siu, M., and Kimball, O. (2017). Improving deliverable speech-to-text systems with multilingual knowledge transfer. In *18th Annual Conference of the International Speech Communication Association (Interspeech)*, pages 127–131.

Manohar, V., Povey, D., and Khudanpur, S. (2017). JHU kaldi system for arabic MGB-3 ASR challenge using diarization, audio-transcript alignment and transfer learning. In *IEEE Automatic Speech Recognition and Understanding Workshop, (ASRU)*, pages 346–352.

Oard, D. (2012). Can automatic speech recognition replace manual transcription? In Doug Boyd, et al., editors, *Oral History in the Digital Age*, Washington, D.C. Institute of Museum and Library Services.

Peddinti, V., Povey, D., and Khudanpur, S. (2015). A time delay neural network architecture for efficient modeling of long temporal contexts. In *16th Annual Conference of the International Speech Communication Association (Interspeech)*, pages 3214–3218.

Picheny, M., Tüske, Z., Kingsbury, B., Audhkhasi, K., Cui, X., and Saon, G. (2019). Challenging the boundaries of speech recognition: The malach corpus. In *To appear in the 20th Annual Conference of the International Speech Communication Association (Interspeech)*.

Povey, D., Ghoshal, A., Boulianne, G., Burget, L., Glembek, O., Goel, N., Hannemann, M., Motlicek, P., Qian, Y., Schwarz, P., Silovsky, J., Stemmer, G., and Vesely, K. (2011). The kaldi speech recognition toolkit. In *IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. IEEE Signal Processing Society, December.

Povey, D., Peddinti, V., Galvez, D., Ghahremani, P., Manohar, V., Na, X., Wang, Y., and Khudanpur, S. (2016). Purely sequence-trained neural networks for ASR based on lattice-free MMI. In *17th Annual Conference of the International Speech Communication Association (Interspeech)*, pages 2751–2755.

Psutka, J., Ircing, P., Psutka, J. V., Radová, V., Byrne, W. J., Hajic, J., Gustman, S., and Ramabhadran, B. (2002). Automatic transcription of czech language oral history in the MALACH project: Resources and initial experiments. In *Text, Speech and Dialogue, 5th International Conference, TSD*, pages 253–260.

Sak, H., Senior, A. W., and Beaufays, F. (2014). Long short-term memory recurrent neural network architectures for large scale acoustic modeling. In *15th Annual Conference of the International Speech Communication Association (Interspeech)*, pages 338–342.

Usbeck, R., Ngomo, A. N., Haarmann, B., Krithara, A., Röder, M., and Napolitano, G. (2017). 7th open challenge on question answering over linked data (QALD-7). In *Semantic Web Challenges - 4th SemWebEval Challenge at (ESWC)*, pages 59–69.

Waibel, A. H., Hanazawa, T., Hinton, G. E., Shikano, K., and Lang, K. J. (1989). Phoneme recognition using time-delay neural networks. *IEEE Trans. Acoustics, Speech, and Signal Processing*, 37(3):328–339.

Wang, D. and Zheng, T. F. (2015). Transfer learning for speech and language processing. In *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA)*, pages 1225–1237.

Xu, H., Su, H., Ni, C., Xiao, X., Huang, H., Chng, E. S., and Li, H. (2016). Semi-supervised and cross-lingual knowledge transfer learnings for DNN hybrid acoustic models under low-resource conditions. In *17th Annual Conference of the International Speech Communication Association (Interspeech)*, pages 1315–1319.

9. Language Resource References

Baum, D., Schneider, D., Bardeli, R., Schwenninger, J., Samlowski, B., Winkler, T., and Köhler, J. (2010). DiSCo - A german evaluation corpus for challenging problems in the broadcast domain. In *International Conference on Language Resources and Evaluation (LREC)*.

Cieri, C., Miller, D., and Walker, K. (2004). The fisher corpus: a resource for the next generations of speech-to-text. In *International Conference on Language Resources and Evaluation (LREC)*.

Godfrey, J. J., Holliman, E. C., and McDaniel, J. (1992). Switchboard: telephone speech corpus for research and development. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP*, volume 1, pages 517–520 vol.1, March.

Panayotov, V., Chen, G., Povey, D., and Khudanpur, S. (2015). Librispeech: An ASR corpus based on public domain audio books. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP*, pages 5206–5210.

Richey, C., Barrios, M. A., Armstrong, Z., Bartels, C., Franco, H., Graciarena, M., Lawson, A., Nandwana, M. K., Stauffer, A. R., van Hout, J., Gamble, P., Hetherly, J., Stephenson, C., and Ni, K. (2018). Voices obscured in complex environmental settings (voices) corpus. In *19th Annual Conference of the International Speech Communication Association (Interspeech)*, pages 1566–1570.

Stadtschnitzer, M., Schwenninger, J., Stein, D., and Köhler, J. (2014). Exploiting the large-scale german broadcast corpus to boost the fraunhofer IAIS speech recognition system. In *International Conference on Language Resources and Evaluation (LREC)*, pages 3887–3890.