



This is a repository copy of *Investigating the impact of cross-lingual acoustic-phonetic similarities on multilingual speech recognition*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/189118/>

Version: Published Version

---

**Proceedings Paper:**

Farooq, M.U. and Hain, T. [orcid.org/0000-0003-0939-3464](https://orcid.org/0000-0003-0939-3464) (2022) Investigating the impact of cross-lingual acoustic-phonetic similarities on multilingual speech recognition. In: Interspeech 2022 - 23rd Annual Conference of the International Speech Communication Association. Interspeech 2022 - Human and Humanizing Speech Technology, 18-22 Sep 2022, Incheon, Korea. International Speech Communication Association , pp. 3849-3853.

<https://doi.org/10.21437/Interspeech.2022-10916>

---

© 2022 ISCA. Reproduced in accordance with the publisher's self-archiving policy.

**Reuse**

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>



# Investigating the Impact of Cross-lingual Acoustic-Phonetic Similarities on Multilingual Speech Recognition

Muhammad Umar Farooq, Thomas Hain

Speech and Hearing Research Group, University of Sheffield, UK.

{mufarooq1,t.hain}@sheffield.ac.uk

## Abstract

Multilingual automatic speech recognition (ASR) systems mostly benefit low resource languages but suffer degradation in performance across several languages relative to their monolingual counterparts. Limited studies have focused on understanding the languages behaviour in the multilingual speech recognition setups. In this paper, a novel data-driven approach is proposed to investigate the cross-lingual acoustic-phonetic similarities. This technique measures the similarities between posterior distributions from various monolingual acoustic models against a target speech signal. Deep neural networks are trained as *mapping networks* to transform the distributions from different acoustic models into a directly comparable form. The analysis observes that the languages ‘closeness’ can not be truly estimated by the volume of overlapping phonemes set. Entropy analysis of the proposed mapping networks exhibits that a language with lesser overlap can be more amenable to cross-lingual transfer, and hence more beneficial in the multilingual setup. Finally, the proposed posterior transformation approach is leveraged to fuse monolingual models for a target language. A relative improvement of  $\sim 8\%$  over monolingual counterpart is achieved.

**Index Terms:** automatic speech recognition, multilingual, acoustic-phonetic similarities, model fusion

## 1. Introduction

Multilingual automatic speech recognition (ASR) systems have attained significant attention over the past decade. Motivation for multilingual speech recognition includes; (i) having a single unified model capable of recognising speech of diverse languages [1, 2] and (ii) using shared language representations to improve ASR performance in the low resource settings [3, 4, 5, 6, 7, 8, 9, 10].

Previous works show that multilingual ASRs do not yield significant reduction in Word Error Rate (WER) for various languages including English, German, French and several others [1, 2, 11]. Many open source data sets are available for these languages and are generally regarded as resource-rich languages. Most of the multilingual models, which include resource-rich languages, are trained on unbalanced data. So, degradation in the performance of these languages is attributed to increased confusion for them in the multilingual setup [1]. *Conneau et al.* [11] mention language interference as the reason of the increased error rate. However, no proof of these concepts could be found in the literature.

Earlier studies, in the context of efficient data sharing among languages for multilingual setup, are on Context Independent (CI) acoustic models [12]. Context Dependent (CD)

acoustic-phonetic similarity was first studied by *Imperl et al.* [13] and later extended by *Le et al.* [14]. Both of these studies measure distance between two polyphones as a weighted sum of monophonic distances of these polyphones. In [14], a knowledge based approach is applied which needs considerable manual efforts as it requires a hierarchical graph to measure the distance between monophones. Furthermore, the phonemes are clustered on IPA symbols.

Recently, some efforts have been made to interpret the learning of multilingual speech recognition systems [15, 16]. Phoneme Error Rate (PER) of each phoneme in monolingual ASR was compared with that of multilingual system [15]. However, no monotonic trend was observed with the growing number of languages the phoneme shared. The authors described this as “*unexpected*” because the phonemes shared by more languages provide more training data and thus the expected error trend would be decreasing. Motivated by the fact that many languages with significant phonemes overlap pose performance degradation in the multilingual setups [1, 2, 11, 15, 16], the objective and contribution of this work is to study acoustic-phonetic similarities to understand if the cross-lingual phoneme sharing is truly a sharing and how does it impact on multilingual speech recognition setups?

To that end, a novel technique is proposed to estimate the cross-lingual acoustic-phonetic similarities for CD hybrid DNN-HMM acoustic models. Hybrid DNN-HMM system is preferred over end-to-end (e2e) modelling to avoid influence of the entangled language model in e2e speech recognition systems. An equal amount of three West Germanic languages (English, German and Dutch) is used to study the impact on multilingual performance. Behaviour of the monolingual acoustic models against a speech signal of the target language is studied by differentiating the posterior distributions. To compare distributions of source and target AMs, a separate regression neural network is trained for each  $\langle source, target \rangle$  pair to map posteriors from a source language AM to the posteriors of the target language AM.

## 2. Cross-lingual acoustic-phonetic similarities

The motivation of research into multilingual speech recognition is based on an assumption that the articulatory representations of phonemes are very close across the languages and can be considered language independent units [17]. However, several languages with substantial cross-lingual phoneme sharing exhibit poorer performance in multilingual setups. This calls for a study to understand the reason of degradation or improvement in multilingual setups when compared with corresponding monolingual systems.

Hybrid DNN-HMM systems yield better performance than the conventional GMM-HMM based ASRs and outperform e2e ASRs with the limited amounts of training data [18]. Further-

This work was partly supported by LivePerson Inc. at the LivePerson Research Centre.




Figure 1: Proposed system architecture

more, the output from e2e speech recognition systems is influenced by the entangled language model which adversely effects the acoustic analysis.

In hybrid speech recognition systems, a deep neural network is trained to produce a posterior distribution of tied states of HMM models. The total number of states (and thus output layer dimension of DNN) is reduced by clustering many poly-phonemes together. Each language yields a different phonetic decision tree in its monolingual ASR. Thus the number of tied states differs for each language and the posterior distributions are not directly comparable across the languages.

In this work a data-driven approach, to transform posteriors from diverse models to a directly comparable form, is proposed. A similarity measure is calculated between the posterior distributions from different models against a given speech signal to estimate the cross-lingual acoustic-phonetic similarities.

### 2.1. Similarity measure

Let  $M_A$  and  $M_{S_i}$  be the monolingual acoustic models of target and source languages respectively. The target language is the language for which the similarity is being measured against the source languages. A regression neural network  $N_{S_iA}$  is trained to translate posteriors  $P_{S_i}$  of dimension  $d_{S_i}$  from  $M_{S_i}$  to the posteriors  $P_{S_iA}$  of dimension  $d_A$  where  $d_A$  is the dimension of posteriors from  $M_A$ . An underlying assumption is that this *mapping network* is able to learn some language related relationships between posterior distributions of source and the target acoustic models. For example, the network could learn the phonemes of target language which are more amenable to cross-lingual transfer than the others. A few hours of speech data can give thousands of examples that provide sufficient training data for *mapping network*. Kullback-Leibler (KL) divergence, the most widely used measure to differentiate two posterior distributions [19], is calculated as a similarity measure between the posterior distributions from the target language AM and the mapped posterior distributions from the source language AM. The proposed system architecture is shown in Fig. 1.

Let  $X = \{x_1, x_2, \dots, x_T\}$  be a set of observations of target language, for which posterior distributions ( $P^Z = \{p_1, p_2, \dots, p_T\}$  where  $Z \in (A, S_iA)$ ) are attained from all monolingual acoustic models. Posteriors from source acoustic models ( $P^{S_i}$ ) are mapped to target posteriors ( $P^{S_iA}$ ) using *mapping network*  $N_{S_iA}$ . The similarity between a source and the target language for a given set of the observations is calculated as

$$D_X(M_T, M_{S_i}) = \frac{\sum_{t=1}^T p_t^A \cdot (\log p_t^A - \log p_t^{S_iA})}{T} \quad (1)$$

## 3. Experimental Setup

### 3.1. Data set

Experiment results are reported using three languages (English *en*, German *de* and Dutch *nl*) of West Germanic family. From previous works, the performance of these languages is either degraded or show a very minor improvement in multilingual

Table 1: % Cross-lingual phoneme shares

		Shares of		
		<i>en</i>	<i>de</i>	<i>nl</i>
shared with	<i>en</i>	100%	70.58%	69.48%
	<i>de</i>	72.96%	100%	88.59%
	<i>nl</i>	69.8%	85.14%	100%

setups despite of a sufficient number of the shared phonemes. The sharing of cross-lingual phonemes is tabulated in Table 1.

The experiments are carried out using portions of the Multilingual LibriSpeech (MLS) data set [20]. For training and evaluation of monolingual AMs, 30 hours and 2 hours are randomly sampled from corresponding sets of the MLS corpus respectively. Limited amount of data for each language is used in this study due to several reasons.

- The amount of data is restricted as an ablation study to the argument that the performance of these languages degrades in the multilingual setups due to their very strong monolingual counterparts (trained on much more data being resource-rich languages) [1].
- On phonetic level, 30 hours are sufficient for acoustic-phonetic similarities analysis as they provide millions of examples for *mapping network* training.
- Experiments with the limited data for model fusion, provides a realistic scenario for low resource languages and the technique can be extended for resource-deficient languages if the outcome is encouraging.

Baseline multilingual ASR is trained by mixing data from all the languages and hence trained on 90 hours of speech data. To train the *mapping network*, the 30 hours train set is further divided into 29 hours of training and 1 hour for validation.

### 3.2. Similarity measure

All acoustic models used for experiments in this work are trained using lattice-free MMI criterion (LF-MMI) [18]. Though the study can easily be extended for polyphones of any context, left biphones are modelled for experimentation here and the term *biphones* will be used during the remainder of the paper.

Given the partially overlapped phoneme sets of two languages (the target and a source language), there are the following subsets of biphones:

- *Shared biphones*: Due to the overlapping cross-lingual phoneme set, there would be some biphones which occur in both (source and the target) languages.
  - *Shared seen biphones (SS)*: The set of shared biphones which are seen by both languages during training of their acoustic models.
  - *Shared unseen biphones (SU)*: The set of shared biphones which are never seen by source language in its train set.
- *Unshared biphones (U)*: The biphones of the target language that are never seen by source language due to non-overlapping phonemes.

A similarity measure is calculated for each of the aforementioned case and an analysis of trends in distance measure is carried out to comprehend the language similarities and their behaviour in the multilingual setups.

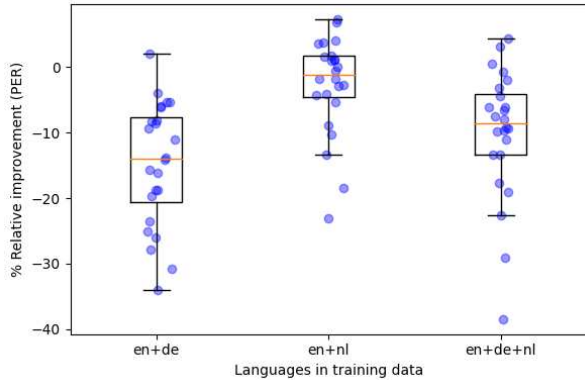


Figure 2: % Relative improvement in PER per shared phoneme compared with monolingual ASR for *en* target language

### 3.3. Model fusion

The *mapping network* is further used to map posterior distributions from several models to a uniform (target language) dimension. The weighted sum of these mapped posteriors is then used for ASR decoding.

For a given observation at time  $t$ , the final posterior vector is given as;

$$p_t^F = w_T \cdot p_t^A + \sum_{i=1}^N w_i \cdot p_t^{S_i A} \quad (2)$$

where  $w$  are the scalar weights assigned to each posterior vector such that  $\sum w = 1$  and  $N$  is the number of source languages.

## 4. Results and Discussion

### 4.1. Baseline multilingual ASR

Monolingual (*mono*) baseline systems are the language dependent acoustic, pronunciation and language models which are trained on a language specific data set. The train sets of all the languages are then mixed to train multilingual (*multi*) acoustic and language models. The multilingual acoustic model is then used with monolingual Language Model (LM) of the target language which is termed as *mono-lm* in the reported results. The results of the baseline systems for all the languages are given in Table 2 in terms of WER and PER. Since only the language model is changed for *mono-lm*, thus the PER remains unchanged when compared with that of multilingual system. The results show that the error for all the languages increases in multilingual setup in spite of the balanced data duration for each language. This indicates that the reason of performance degradation in multilingual setups can not only be attributed to rich resources of these languages or unbalanced data sampling.

According to the assumption of multilingual systems [17] discussed earlier in section 2, if the articulatory representations of phonemes are considered language independent units then the performance of shared phonemes should improve with the

Table 2: Baseline ASR performance

Language	WER/PER		
	<i>mono</i>	<i>multi</i>	<i>mono-lm</i>
English ( <i>en</i> )	43.84/28.10	47.48/31.06	46.40/31.06
German ( <i>de</i> )	37.77/26.86	40.81/28.35	38.11/28.35
Dutch ( <i>nl</i> )	37.94/21.40	58.84/36.16	52.33/36.16

Table 3: Posterior distribution similarity for the *en* test set

AM	Biphone subsets	% Correct SAMC	KL-Div (SAMC)	Entropy (SAMC)
<i>de</i>	<i>SS</i>	37.08	1.32 (0.32)	2.43 (1.4)
	<i>RSS</i>	13.30	1.75 (0.84)	3.22 (2.64)
	<i>RSU</i>	-	1.88	3.40
	<i>RU</i>	-	2.27	3.59
<i>nl</i>	<i>SS</i>	38.04	1.23 (0.32)	2.21 (1.29)
	<i>RSS</i>	12.90	1.64 (0.82)	2.95 (2.44)
	<i>RSU</i>	-	1.71	2.86
	<i>RU</i>	-	1.81	3.1

more training data in the multilingual systems. The languages, being studied here, have an overlapping set of 24 phonemes. As a case study of *en* as the target language, the relative improvement in PER of shared phonemes is analysed with gradually increasing the languages in the training data. It is evident from the Fig. 2 that even the performance of shared phonemes is degraded in bilingual and multilingual setups, though *nl* is less detrimental than *de* for the *en* language.

### 4.2. Similarity analysis

Table 3 shows the similarity measure for English (*en*) test set when passed through German (*de*) and Dutch (*nl*) acoustic models. In clustering, shared seen biphones (*SS*) may share the same cluster with unseen (*SU*) or unshared biphones (*U*) and vice versa. So for insightful observations, analysis is restricted to the clusters which have only one biphone (shown with “*R*” prefix). It is clear from the results that KL divergence increases from shared seen biphones towards unshared biphones.

For shared seen biphone sets (*SS* and *RSS*), being present in both (source and the target) languages, mapping network should learn one-to-one mapping to same biphone class if the multilingual assumption holds true. This measure is calculated as percentage of these biphones recognised correctly by source acoustic model and reported as ‘*correct Source Acoustic Model Class (SAMC)*’ in the results. The results (low percentage of “*Correct SAMC*” in tables) show that the source acoustic models are not very good at recognising these biphones but the lower values of KL divergence (*KL-Div* in parenthesis) indicate that it is easier for *mapping network* to learn a one-to-one mapping in these cases. It implies that the source ASR has a pattern in errors of these biphones sets which mapping network could learn easily. On analysis, it appears that the source acoustic models confuse these biphones with several close biphones. For example, a biphone  $\setminus i \setminus, \setminus z \setminus$  from English test set is frequently confused with  $\setminus i \setminus, \setminus s \setminus$  and  $\setminus \epsilon \setminus, \setminus s \setminus$  by German acoustic model.

To study the confidence of the *mapping network* in the

Table 4: Posterior distribution similarity for the *nl* test set

AM	Biphones subset	% Correct SAMC	KL-Div (SAMC)	Entropy (SAMC)
<i>en</i>	<i>SS</i>	46.51	1.33 (0.37)	1.75 (0.95)
	<i>RSS</i>	15.43	1.87 (1.26)	2.36 (1.77)
	<i>RSU</i>	-	2.113	2.66
	<i>RU</i>	-	2.173	2.69
<i>de</i>	<i>SS</i>	40.90	1.44 (0.46)	1.89 (1.19)
	<i>RSS</i>	17.30	1.88 (1.15)	2.33 (1.81)
	<i>RSU</i>	-	2.04	2.43
	<i>RU</i>	-	2.13	2.58

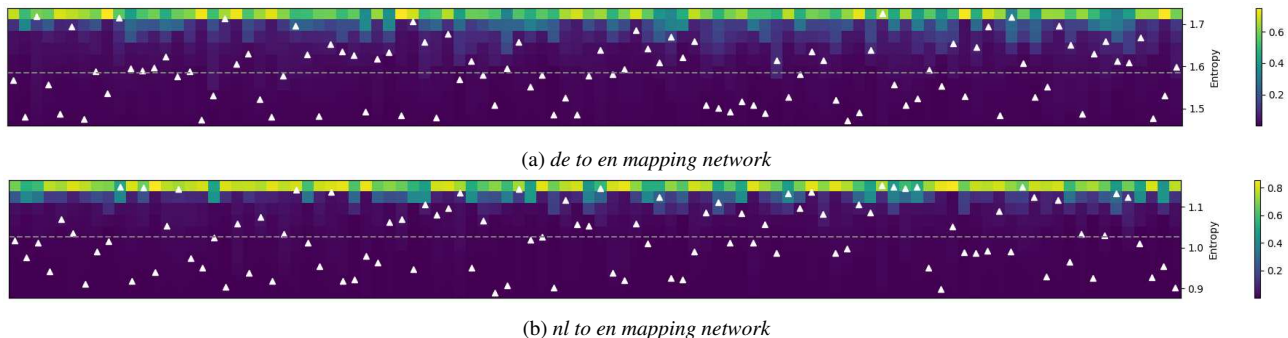


Figure 3: Posteriorgram and entropy plot of  $N_{de-en}$  and  $N_{nl-en}$  mapping networks with one-hot vectors as input. Behaviour for only hundred source biphone classes is shown. Posteriorgram shows sorted probabilities of top ten mapped classes. Each box on (i) horizontal axis is a one-hot vector of source language and on (ii) vertical axis is the probability of a mapped output class

Table 5: Posterior distribution similarity for the *de* test set

AM	Biphones subset	% Correct SAMC	KL-Div (SAMC)	Entropy (SAMC)
<i>en</i>	<i>SS</i>	43.56	0.83 (0.22)	1.76 (1.07)
	<i>RSS</i>	14.90	1.21 (0.74)	2.52 (2.12)
	<i>RSU</i>	-	1.15	2.21
	<i>RU</i>	-	1.27	2.54
<i>nl</i>	<i>SS</i>	36.37	1.05 (0.31)	1.95 (1.24)
	<i>RSS</i>	13.68	1.38 (0.82)	2.5 (2.09)
	<i>RSU</i>	-	1.54	2.67
	<i>RU</i>	-	1.41	2.56

mapped classes, it is sequentially fed the one-hot vectors as input. As a case study, the posteriorgram and entropy of *mapping networks* from source languages to *en* is analysed. Behaviour for the top  $n$  source classes, which are mapped to the target classes with minimum entropy (more confidently), is observed. Entropy and posteriorgram from both networks,  $N_{de-en}$  and  $N_{nl-en}$ , for same  $n$  ( $n = 100$ ) is visualised in Fig. 3. Only top ten most probable mapped output classes are shown in the sorted posteriorgram. It can be seen that for the same value of  $n$ , the entropy range is lower for *nl-en* mapping network than the range for *de-en* network. It evidences that *nl-en* mapping network could learn better mappings and *nl* phonemes are more amenable to transfer to *en* phonemes. Furthermore, the entropy for each biphones subset is also measured which also shows the same trend as the KL divergence. It implies that the language similarities can be estimated through the entropy of the mapping network only without measuring KL-div between target and mapped posteriors.

In the case of unseen (*RSU*) and unshared (*RU*) biphones, the source model tries to map them to the nearest biphone clusters. However, the performance of unshared biphones mapping goes further down in most of the cases. Almost the same pattern is observed for the remaining two languages (Table 4 and 5). The average cross-lingual KL divergence along with the WER of bilingual ASRs is tabulated in Table 6. Diagonal entries are

Table 6: Mean KL-divergence (as the cross-lingual similarity measure) and bilingual ASR performance

Target Language	Source Languages(KL-Div/% WER)		
	en	de	nl
en	0/43.84	1.56/46.35	<b>1.44/44.46</b>
de	<b>1.00/38.94</b>	0/37.77	1.18/39.46
nl	<b>1.60/42.32</b>	1.61/44.03	0/37.94

Table 7: Performance of model fusion in % WER

Language	mono	mono-lm	mf	Rel. imp
English ( <i>en</i> )	43.84	47.48	<b>43.43</b>	0.94%
German ( <i>de</i> )	37.77	40.81	<b>36.16</b>	4.26%
Dutch ( <i>nl</i> )	37.94	58.84	<b>34.94</b>	7.91%

the monolingual ASRs. These results can be seen in comparison with cross-lingual phoneme sharing of Table 1. For example, Table 1 infers that *de* shares more phonemes with *nl* compared to *en*, but the mean KL divergence is smaller for *en* (in Table 6). It implies that *en* is more closer to *de* than *nl* and this claim is corroborated by % WER of bilingual ASRs in Table 6. So, the phoneme sharing statistics are not a very informative metric to measure languages closeness and thus lead to degradation in the multilingual setups.

### 4.3. Model fusion for multilingual ASR

A multilingual acoustic model is imitated by fusing the target language and the mapped source language posteriors. The fusion is the linear weighted sum of all of these posterior distributions. In Table 7, the results of proposed model fusion (*mf*) approach is compared with *mono* and *mono-lm* of Table 2. The setup of fusion technique makes results directly comparable to those of *mono-lm* but the fusion outperforms even the *mono* ASR. So, the reported relative improvement here is in comparison with the *mono* ASR. However, fused posteriors give a relative gain of 6.4% to 30% when compared with *mono-lm*.

## 5. Conclusion

In this work, cross-lingual acoustic-phonetic similarities are estimated by comparing the posterior distributions from source and the target acoustic models. A regression neural network is trained to map source languages posteriors to the target language posteriors. This study reveals that the behaviour of different languages for multilingual ASRs is more complex than predicting from cross-lingual phoneme sharing perspective. The languages which share more phonemes, does not guarantee performance gain in multilingual setups. The analysis observes that the phonemes with identical representations across languages are not acoustically identical. Finally, the mapped posteriors are fused for decoding of a target language. A maximum gain of  $\sim 8\%$  in relative improvement over the monolingual system is achieved. The relative gain increases to 30% when compared with the fusion model's counterpart.

## 6. References

- [1] V. Pratap, A. Sriram, P. Tomasello, A. Hannun, V. Liptchinsky, G. Synnaeve, and R. Collobert, “Massively Multilingual ASR: 50 Languages, 1 Model, 1 Billion Parameters,” in *Proc. Interspeech 2020*, 2020, pp. 4751–4755.
- [2] W. Hou, Y. Dong, B. Zhuang, L. Yang, J. Shi, and T. Shinzaki, “Large-Scale End-to-End Multilingual Speech Recognition and Language Identification with Multi-Task Learning,” in *Proc. Interspeech 2020*, 2020, pp. 1037–1041.
- [3] S. T. Abate, M. Y. Tachbelie, and T. Schultz, “Multilingual acoustic and language modeling for ethio-semitic languages,” in *Proc. Interspeech 2020*, 2020, pp. 1047–1051.
- [4] M. Y. Tachbelie, S. T. Abate, and T. Schultz, “Development of multilingual asr using globalphone for less-resourced languages: The case of ethiopian languages,” in *Proc. Interspeech 2020*, 2020, pp. 1032–1036.
- [5] D. Imseng, P. Motlicek, H. Bourlard, and P. Garner, “Using out-of-language data to improve an under-resourced speech recognizer,” *Speech Communication*, vol. 56, p. 142–151, 01 2014.
- [6] M. Karafiát, M. K. Baskar, P. Matějka, K. Veselý, F. Grézl, and J. Černocký, “Multilingual blstm and speaker-specific vector adaptation in 2016 but babel system,” in *IEEE SLT*, 2016, pp. 637–643.
- [7] L. Besacier, E. Barnard, A. Karpov, and T. Schultz, “Automatic speech recognition for under-resourced languages: A survey,” *Speech Communication*, vol. 56, pp. 85–100, 2014.
- [8] N. T. Vu and T. Schultz, “Multilingual multilayer perceptron for rapid language adaptation between and across language families,” in *Proc. Interspeech 2013*, 2013, pp. 515–519.
- [9] S. Tong, P. N. Garner, and H. Bourlard, “Cross-lingual adaptation of a ctc-based multilingual acoustic model,” *Speech Communication*, vol. 104, pp. 39–46, 2018.
- [10] J.-T. Huang, J. Li, D. Yu, L. Deng, and Y. Gong, “Cross-language knowledge transfer using multilingual deep neural network with shared hidden layers,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 7304–7308.
- [11] A. Conneau, A. Baeovski, R. Collobert, A. Mohamed, and M. Auli, “Unsupervised Cross-Lingual Representation Learning for Speech Recognition,” in *Proc. Interspeech 2021*, 2021, pp. 2426–2430.
- [12] J. Kohler, “Multi-lingual phoneme recognition exploiting acoustic-phonetic similarities of sounds,” in *Proceeding of Fourth International Conference on Spoken Language Processing. ICSLP ’96*, vol. 4, 1996, pp. 2195–2198 vol.4.
- [13] B. Imperl, Z. Kacic, B. Horvat, and A. Zgank, “Agglomerative vs. tree-based clustering for the definition of multilingual set of triphones,” in *2000 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No.00CH37100)*, vol. 3, 2000, pp. 1273–1276 vol.3.
- [14] V. B. Le, L. Besacier, and T. Schultz, “Acoustic-phonetic unit similarities for context dependent acoustic model portability,” in *2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*, vol. 1, 2006.
- [15] P. Želasko, L. Moro-Velázquez, M. Hasegawa-Johnson, O. Scharenborg, and N. Dehak, “That Sounds Familiar: An Analysis of Phonetic Representations Transfer Across Languages,” in *Proc. Interspeech 2020*, 2020, pp. 3705–3709.
- [16] S. Feng, P. Želasko, L. Moro-Velázquez, A. Abavisani, M. Hasegawa-Johnson, O. Scharenborg, and N. Dehak, “How phonotactics affect multilingual and zero-shot asr performance,” in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 7238–7242.
- [17] T. Schultz and A. Waibel, “Language-independent and language-adaptive acoustic modeling for speech recognition,” *Speech Commun.*, vol. 35, no. 1–2, p. 31–51, 2001.
- [18] D. Povey, V. Peddinti, D. Galvez, P. Ghahremani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, “Purely Sequence-Trained Neural Networks for ASR Based on Lattice-Free MMI,” in *Proc. Interspeech 2016*, 2016, pp. 2751–2755.
- [19] S. Kullback and R. A. Leibler, “On Information and Sufficiency,” *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79 – 86, 1951.
- [20] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, and R. Collobert, “Mls: A large-scale multilingual dataset for speech research,” *ArXiv*, vol. abs/2012.03411, 2020.