

# A Survey of Multilingual Models for Automatic Speech Recognition

Hemant Yadav, Sunayana Sitaram

Microsoft Research India

Bangalore

t-hyadav@microsoft.com, sunayana.sitaram@microsoft.com

## Abstract

Although Automatic Speech Recognition (ASR) systems have achieved human-like performance for a few languages, the majority of the world’s languages do not have usable systems due to the lack of large speech datasets to train these models. Cross-lingual transfer is an attractive solution to this problem, because low-resource languages can potentially benefit from higher-resource languages either through transfer learning, or being jointly trained in the same multilingual model. The problem of cross-lingual transfer has been well studied in ASR, however, recent advances in Self Supervised Learning are opening up avenues for unlabeled speech data to be used in multilingual ASR models, which can pave the way for improved performance on low-resource languages. In this paper, we survey the state of the art in multilingual ASR models that are built with cross-lingual transfer in mind. We present best practices for building multilingual models from research across diverse languages and techniques, discuss open questions and provide recommendations for future work.

**Keywords:** speech recognition, multilingual, low-resource languages

## 1. Introduction

The field of Automatic Speech Recognition (ASR) has made significant progress over the last few years due to the advent of Deep Neural Network (DNN) based models. Most notably, on benchmarks such as Switchboard, models have now achieved human parity (Xiong et al., 2016), which means that models are able to transcribe human speech with the same or better accuracy as humans. However, due to the large data requirement of DNN-based models, most of these advances have been restricted to a few languages of the world, which have large datasets of transcribed speech. Languages that do not have much transcribed and/or untranscribed speech, that is, low-resource languages have been left behind.

(Joshi et al., 2020) classify languages based on resources for Natural Language Processing (NLP) and find that almost 90% of the world’s population speaks languages that are very low resource and are thus not benefited by language technologies. Since speech datasets are usually more expensive to build than text datasets, the situation may be even more stark for ASR. This is a lost opportunity for serving language communities where speech is the only medium of expression, owing to low literacy or the lack of a standardized writing system.

The problem of ASR for low resource languages has been well studied starting from traditional HMM-GMM based models, to the transformer-based models used today. The main technique used for addressing the lack of resources in the target language is cross-lingual transfer, in which models or resources from a higher resource language are adapted to the target low resource language using adaptation data. Cross-lingual transfer may take place using models trained on a single high resource language, or multiple languages with varying amounts of resources. The intuition behind these tech-

niques is that the lower resource language benefits from the language invariant features learned by the model from the large amount of data in the high resource language.

Recently, there has been a paradigm shift in NLP research with the advent of pre-trained models such as BERT (Devlin et al., 2018). These models first use unlabeled data (raw text) for pre-training and are then fine-tuned for a downstream task using labeled task-specific data. The intuition behind this approach is that the model can learn features from unlabeled data that can be shared across tasks. Massively Multilingual Language Models (MMLMs) (Devlin et al., 2019; Conneau et al., 2020b) are pre-trained on a large amount of unlabeled text from multiple languages (around 100) and then fine-tuned for a particular language and task. MMLMs have shown impressive cross-lingual transfer (Wu and Dredze, 2019), leading to performance gains on languages that have no labeled data (through zero-shot learning) or a small amount of labeled data (few-shot learning). Recently, the idea of using unlabeled data for pre-training has also been used in the field of speech processing (Chung et al., 2021), which could greatly solve the challenges of low-resource ASR if successful, since there may be large amounts of untranscribed speech available in many languages. An additional benefit of multilingual models is that training and maintenance costs can be lowered by reusing the same model for multiple languages, which can benefit low-resource languages having limited budgets.

In this paper, we survey the progress made in multilingual ASR. For the purposes of this paper, we define multilingual ASR models as models that are trained with the aim of aiding cross-lingual transfer to eventually benefit other (usually low-resource) languages. We classify them as follows (i) ASR models trained us-

ing labelled data in one or multiple languages (ii) ASR models that are first pre-trained using unlabeled data from one or multiple languages and then fine-tuned using labeled data. In our survey, we set out to answer the following questions:

- Are multilingual models better in terms of performance compared to monolingual models for high and low-resource languages?
- Can multilingual ASR models exploit unlabeled data for improving performance on low-resource languages?
- What are the best practices for building multilingual ASR models?
- Where should resources be invested for improving low-resource ASR in the era of multilingual models?
- What are the open questions and main challenges that need to be addressed going forward?

The paper is organized as follows. Section 2 describes the general ASR training pipeline and resources required for building ASR models. In Section 3, we survey multilingual models for ASR categorized by whether or not they use unlabeled data. In Section 4, we list the key findings and open questions that still need to be addressed. Section 5 concludes.

## 2. ASR training and resources

ASR is the task of converting a spoken utterance into a sequence of words. It can be broken down into three broad steps.

(i) The first step is to convert the raw audio into a compact feature vector, for example MFCC, PLP, Spectrogram, *etc.* (ii) The next step is to learn an acoustic model, usually referred to as an encoder to associate these features to higher level speech or orthographic units (phonemes or graphemes). This is achieved either using Gaussian Mixture Models (GMM) (Khan et al., 2018) or a Deep Neural Network (DNN) (Hannun et al., 2014; Amodei et al., 2016; Li et al., 2019) with HMMs, where the GMM/DNN acts as a function approximator and the HMM adds context capabilities to the encoder. With advances in deep learning, Long Short Term Memory networks (LSTMs) have replaced both the GMM/DNN and HMM models, since LSTMs act as function approximators and also learn context-dependence. The two most popular training frameworks used for acoustic modeling are Expectation-maximization (EM) (Dempster et al., 1977) and Connectionist Temporal Classification (CTC) (Graves et al., 2006) respectively. Using the CTC loss function GMM/DNN-HMM system can be replaced by powerful auto-regressive models like LSTMs and Transformers. (iii) Lastly, the output speech units are conditioned on a pre-trained language model (LM) (Heafield, 2011;

Devlin et al., 2018) such that the output sentence has a low perplexity.

Traditional techniques optimize the ASR task in three steps independently as described above. Both GMM-based and early DNN-based techniques relied on the following language resources (i) A large corpus of speech transcribed by native speakers for training, (ii) a pronunciation lexicon mapping words in the language to phonemes, (iii) a phone set for the language, (iv) a large corpus of text for the language model (v) transcribed data for evaluation.

With the the availability of large annotated datasets for some languages, current state-of-the-art (SOTA) techniques solve the task in an end-to-end fashion (Gulati et al., 2020; Synnaeve et al., 2019; Han et al., 2020) via supervised learning using only transcribed speech, in which sounds are directly mapped to graphemes/sub-word-units present in the language (Zhou et al., 2021). This forgoes the need for a lexicon and language model, which are learned implicitly by the model. Although end-to-end training seems attractive for low-resource languages where lexicons and text data may not be available, these models typically require larger amounts of labeled speech data than their traditional counterparts to reach the same level of performance (Kannan et al., 2019; Toshniwal et al., 2018).

Recently, self supervised learning (SSL) has been used to pre-train an encoder on unlabelled data using contrastive loss. In this setup, the model learns a general high-level contextual representation of the input data which can potentially be used for any downstream task. The training process is divided into two steps: (i) learning an encoder which maps the raw audio to a high level compact representation, usually done using Convolutional Neural Networks (CNNs) (ii) Reconstruct the future frames given the high level features, using a strong auto-regressive model such as Transformers. Experiments by (Baeviski et al., 2020) have shown that quantizing the high level features before computing the contrastive loss leads to performance gains. These models have achieved SOTA performance in low resource scenarios and have been shown to perform on par, if not better, in high resource settings (Baeviski et al., 2020; Hsu et al., 2021; Chung et al., 2021).

## 3. Multilingual ASR models

In this section, we survey recent research on multilingual ASR models. We categorize multilingual ASR models based on the data that they can be trained on - labeled data, in which each word that is spoken is carefully transcribed by a native speaker, and unlabeled data, which consists of speech without any transcriptions. Labeled data is expensive to create and may not always be feasible to collect if access to native speakers is challenging. Unlabeled data may be available 'in the wild' for many languages, in the form of user-generated content, such as on YouTube.

### 3.1. Models that use only labeled data

#### 3.1.1. Acoustic models

(Heigold et al., 2013) presents an extensive study comparing three major approaches towards building multilingual ASR using the GMM or DNN (or combined) and HMM setup.

- **Language independent feature extraction:** In this approach, feature extraction is learned independent of the languages, so the feature extractor can be used for any language, after which an ASR system can be built on top of it. A major drawback of such a system is that it is not optimised jointly.
- **Transfer learning:** In this approach, a model's learnable weights are initialised from a model trained on a high-resource language, followed by a separate fine-tuning step.
- **Multi-task learning:** In this setting, one unified model is optimized/trained to perform  $n$  tasks, without any separate fine-tuning step. In our case,  $n$  is the number of different languages *i.e.*, and each language is treated as a separate task.

(Heigold et al., 2013) conduct experiments on 11 languages with a total duration of around 10000 hours for training and 35 hours for testing. The authors first compare, on all the 11 languages, GMM-HMM and DNN-HMM based monolingual systems and observe that the later performs better. Therefore they use DNN-HMM architecture for cross-lingual and multilingual experiments. (i) In the cross-lingual experiments, the DNN-HMM model is first trained on 3000 hours of English and then fine-tuned in two settings, either fine-tuning the top two layers or all layers for each language. Fine-tuning the top two layers avoids over-fitting for low resource languages. (ii) In the multilingual experiments, the DNN-HMM model is trained on combined data of all the 11 languages. It performs better than the cross-lingual setup consistently on all the languages. This study shows that **multilingual models trained by combining data of all languages performs better than fine-tuning an English model on the target languages.**

(Thomas et al., 2010) proposes using trained Multi-Layer Perceptron (MLP) features in cross-lingual settings *i.e.*, adapting the learned MLP to a new language using phone mapping without re-initializing the last layer. The authors train an MLP on a combination of German and Spanish and after the phone mapping fine-tune the MLP on 1 hour English dataset, which they treat as the low-resource language. (Thomas et al., 2012) proposes training a four-layered MLP for feature extraction on a German and Spanish followed by fine-tuning on English. In this setup, the first three layers are shared among different languages and the last layer is language specific. This results in a substantial (30%) performance gain compared to the baseline

monolingual GMM-HMM models. This work shows that **feature extraction and phone-mapping followed by fine-tuning of the target language results in significant gains over the monolingual model**, possibly because the source and target languages are from the same language family.

(Tüske et al., 2013a) compare systems built using Bottleneck (BN) features and MFCC features on datasets consisting of 150 hours of speech in German, Spanish, and English. They experiment with two setups (i) cross-lingual: fine-tuning the pre-trained MLP model from one language to a target language. They report a slight performance gain using BN features over monolingual baselines. (ii) multilingual: train one MLP model with all the languages combined. Their experiments show consistent improvement on all the three languages using BN features compared to MFCC features.

(Tüske et al., 2013b) perform extensive experimentation to study the scalability of BN features using four languages *i.e.*, French (317 hours), English (232 hours), German (142 hours), and Polish (110 hours) and train five and nine layer MLP models. The experiments show that BN features trained in multilingual settings perform better compared to BN features trained on the target language only. Lastly, jointly training a language dependent hidden layer with multilingual BN features performs best. This shows that **the choice of features may play an important role in cross-lingual transfer, and that training features on all languages performs better than training only on the target language.**

(Ghoshal et al., 2013) trains the DNN directly, in the hybrid DNN-HMM setup, in sequential order and re-initialize the softmax layer for each language *i.e.*, they train on one language then fine-tune on the second language, followed by fine-tuning on the third language and so on. They experiment with seven European languages in their setup with a total duration of 140hours for training and 14 hours for testing. Their setup shows improvement over the monolingual GMM-HMM and DNN-HMM baselines, however it is not clear whether the improvement is due to sequential training or the additional languages and training data added in each step.

(Huang et al., 2013) train a shared-hidden-layer multilingual DNN (SHL-MDNN), setup on four languages *i.e.*, French (138 hours), German (195 hours), Spanish (63 hours), and Italian (63 hours). Their experiments show that multilingual training improves over monolingual training on all languages. Furthermore they also train models on two unseen languages, English and Chinese, using the trained DNN as a feature extractor. Cross-lingual transfer performs better than training models for English and Chinese from scratch, which shows the learned features are transferable to unseen languages also. It is interesting to note that **the transfer performs well for Chinese too, which is not in the same language family as the languages used to train**

**the feature extractor.**

(Graves et al., 2013) shows that LSTM is a better choice of architecture than DNNs for ASR. Based on this, (Zhou et al., 2017) study LSTM in the shared-hidden-layer multilingual LSTM (SHL-MLSTM) setting and compare it to SHL-MDNN. They show that monolingual SHL-MLSTM performs comparable to multilingual SHL-MDNN and multilingual SHL-MLSTM consistently outperforms multilingual SHL-MDNN on all the seven languages. SHL-MLSTM with residual learning helps them to train deeper networks, which results in even better performance. They hypothesize that residual learning helps overcome the gradient vanishing problem, indicating that **the choice of architecture is important while building larger, deeper models trained on multiple languages**. They also find that **the performance increases as the number of languages in the multilingual model increase**.

The Interspeech 2018 ASR challenge (Srivastava et al., 2018) compared models built for three Indian languages (Gujarati, Tamil and Telugu), in which baselines were monolingual models built using HMM-GMM, DNN-HMM and TDNN techniques. Most systems that improved over the baselines used multilingual training, that is combining data from all languages to build a single model and decoding using language specific Language Models (Billa, 2018). (Sailor and Hain, 2020) work on the same dataset and train four models (i) monolingual, (ii) multilingual, (iii) MTL with Language ID, and (iv) MTL with phoneme recognition. Their experiments show that using language ID as an auxiliary task degrades the performance of multilingual ASR system and propose that phoneme recognition is better choice. Adding **phoneme recognition as an auxiliary task gives small performance gains compared to the multilingual model**.

One other approach is to replace the GMM-HMM and train acoustic models (LSTMs) using the CTC loss function. (Tong et al., 2017) studies training on multiple languages using the CTC loss function *i.e.*, for English (81 hours), French (120 hours) and German (136 hours). In their experiments **training on all the languages combined performs worse compared to monolingual models, and fine-tuning multilingual models on target languages performs better than monolingual models**. They also observe performance gains when the trained multilingual model is fine-tuned on two low-resource unseen European languages. (Müller et al., 2017) shows similar results that feeding language vectors during training performs better than training multilingual models by just combining data.

### 3.1.2. seq2seq models

(Watanabe et al., 2017) used the hybrid attention/CTC framework to study multilingual training on 10 languages, with 1327 hours of speech for training and 150 hours for testing. The authors find that **the performance of some high resource languages degrades**

**in multilingual settings** compared to the monolingual baselines, due to unbalanced training data. To overcome the imbalance, three techniques are proposed in the literature *i.e.*, (i) Appending language ID vector in the input (Toshniwal et al., 2018), which performs better than the baseline, (ii) Sampling, and (iii) combining both, results in best performing models. So, **the degradation in performance for high-resource languages can be overcome by using sampling and language ID information**.

(Toshniwal et al., 2018) studies multilingual training on 9 Indic languages, with around 1500 hours for training and 90 hours for testing. In their work, the authors experiment with four different models using Listen, Attend and Spell (LAS) (Chan et al., 2015) as the architecture. (i) monolingual (baseline) models for each language, (ii) A multilingual/joint model combining all the 9 languages, (iii.) A single model for ASR and predicting the language ID (iv.) A multilingual model conditioned on the language ID as an input. The language ID conditioned model gave the best performance, showing that **including language ID information helps training multilingual ASR systems**. They hypothesize that this is because the model reserves different parts of the network for different languages, while also being able to share language invariant features.

(Zhou et al., 2018) shows similar findings using the transformer based architecture with experiments on six languages from the CALLHOME corpora. Their extensive experimentation with classical SOTA techniques shows that transformer based models perform better than *i.e.*, Mono-DNN, Mono-LSTM, SHL-MDNN, SHL-MLSTM, and SHL-MLSTM-residual models. (Shetty and NJ, 2020) also use transformer based models to show performance gains when using language ID. They conduct experiments in three Indian languages *i.e.*, Gujarati, Tamil, and Telugu. Lastly, they show that retraining the multilingual model on target languages further improves the performance compared to the universal multilingual model.

(Kannan et al., 2019) learn an adapter module after each layer for each language separately for streaming E2E ASR. Similarly to (Toshniwal et al., 2018), they use nine Indic languages for training and testing. The authors train a system combining (i) multilingual RNN-T, (ii) language ID, and (iii) adapter modules. Both (ii) and (iii) performs best on low resource and on par in high resource languages. Adapter modules increase the overall model size compared to language ID training, however **there is a small performance boost for both low and high resource languages when an adapter module is learned for each language separately**.

The multilingual models surveyed so far are trained with fewer than 10 languages, in contrast to the massively multilingual models in NLP that are trained on 100+ languages. Recently (Pratap et al., 2020) studies training of models of size up to 1 billion parameters on

51 languages, the largest till date. The total duration of speech data is around 16000 hours for training and 1000 hours for testing. The authors experiment with (i) a multilingual/joint model combining all the 51 languages, (ii.) a multilingual model conditioned on the language ID as an input and, (iii) multi-head model with 6-language-clusters ie. 6-different-decoders based on language family groupings. The multi-head model gives the best performance because **training the same decoder for multiple languages is beneficial if the languages are similar.**

### 3.2. Models that use unlabeled data

Pre-training can be done on both labelled and unlabelled speech. Pre-training on unlabelled speech data is more practical for low-resource languages because it is easier and less expensive to find unlabeled speech data.

(Conneau et al., 2020a) studied self-supervised multilingual training on 53 languages (56,000 hours) using the XLSR model. The authors train two variants of Wav2Vec2.0 (Baevski et al., 2020) style models (base and large), and perform a series of experiments to study the effectiveness of multilingual SSL training. (i) Pre-training (pre1) (on English only), which outperforms monolingual models on low resource languages. (ii) Multilingual pre-training (pre2) on 10 languages (800 hours), with the same amount of data as in pre1, which outperforms pre1. This shows that **pre-training on multiple languages is a better choice rather than just adding more data from one language.** In both the cases, performance degrades on high resource languages. To address this degradation, they pretrain a larger model with 24 transformer blocks on 53 languages (56,000 hours) and this model shows improvements in high resource languages, in comparison to the pre2 model. This indicates that **SSL may need larger model capacity to show improvements for high resource languages.** (iii) The Pre2 learnt representation transfers well to unseen low resource languages in low resource (50 hours) (iv) Fine-tuning one model for all languages works a little worse on the base model and works slightly better on the large model compared to a monolingual model for each language. This again indicates that models require large capacity to show performance gains on larger datasets and a higher number of languages. (v) Finally, they show that **pre-trained models transfer well if the target language is similar to languages used for pre-training.**

(Gupta et al., 2021) extends these findings by training a model on Indic languages. They use around 10,000 hours for pre-training out of which the majority of the data is in Hindi (50%). Their experiments results are similar to (Conneau et al., 2020a), that **multilingual pre-training outperforms monolingual training**, even in the case of Hindi, which is a large dataset. (Javed et al., 2021) create a speech corpus of 17000 hours of unlabeled data in 40 Indian languages, includ-

ing several very low resource languages. They use this data for pre-training multilingual ASR models and fine-tune the models on 9 Indian languages that have labeled data, leading to SOTA results on multiple benchmarks. Their findings suggest the following: (i) fine-tuning can be performed either monolingually, that is one-by-one for each target language, or combining all languages together. **Jointly fine-tuning a single model with all languages performs as well as multiple models fine-tuned separately.** 2. **The size of pre-training and fine-tuning datasets affect the accuracy of the final model** - contrary to previous findings ((Baevski et al., 2020)), fine-tuning with just ten minutes of speech does not lead to high performance. 3. Building on (Conneau et al., 2020a), **pre-training with a diverse set of languages improves performance on languages not present in the pre-training data**, which indicates that multilingual models may be useful for building speech technologies for very low resource languages, provided that some data for fine-tuning and evaluation is collected. Although these experiments do not include comparisons with monolingual models, the models beat previous SOTA models in terms of performance, which in turn are better performing than monolingual models (Srivastava et al., 2018).

(Wang et al., 2021) proposes to combine supervised learning with labeled data and the CTC loss and SSL with the WAV2VEC2.0 loss in a Multitask Learning (MTL) scheme which they call UniSpeech. This helps the model to learn representations which are useful for speech recognition *e.g.*, phoneme identities or invariance to background noise and accents. Furthermore they randomly replace the contextual representation with their quantized counterparts from the encoder when calculating the CTC loss. According to them this helps aligning the codewords from codebooks to ASR-specific features. The authors perform three sets of experiments based on one or multiple languages used during the pre-training and fine-tuning phase. (i) one-to-one (ii) many-to-one (iii) many-to-many. They report performance gains when using UniSpeech over XLSR (Conneau et al., 2020a) and CTC based model on all the three settings, showing that **the MTL setup is useful for cross-lingual transfer.** One downside of this learning scheme is that the learned features do not transfer to other tasks *e.g.*, emotion or speaker recognition, as the features are specific to the ASR task.

(Billa, 2021) compares ASR systems for three languages, Farsi, Kazakh and Lithuanian built using the following configurations: (i) in-domain labeled training data, which acts as the baseline (ii) mined data from Youtube in each language, followed by fine-tuning with in-domain data (iii) pooling data from 1 and 2 followed by fine-tuning (iv) transfer from a high-resource Arabic model by using the labeled data in each language. They find that fine-tuning either the model with in-domain or pooled data improves over the baseline, showing that **even unlabeled out-of-domain data can be used to**

**improve performance over models built with only labeled in-domain data.** They also show that the fine-tuning approaches are comparable to transfer learning from a high-resource language.

Another technique for using unlabeled data is pseudo labeling, in which untranscribed data is transcribed using an existing ASR system. This is followed by a step to remove noisy labels, after which the transcribed data is added back to train the ASR model, leading to more data to train the model without investing resources for manual transcription. (Javed et al., 2021) shows that pseudo labelling for ASR (Kahn et al., 2020) does not improve performance over 10 Indic languages compared to models built using only transcribed speech. They hypothesize that the generated pseudo labels are noisy due to the difference in distribution between the labeled and unlabeled speech, which is a common scenario in the real world. SSL may be a better technique for addressing domain and distribution differences in labeled and unlabeled data. Another drawback of the pseudolabel technique compared to SSL is that it is designed to be used within the same language and does not transfer.

(Zhang et al., 2021) addresses the lack of transfer across languages by using self-training to learn frame-level representations rather than generating the pseudo labels and call the algorithm Cross-lingual Self-training (XLST). XLST assumes that frame-level acoustic representations could be shared in some degree across different languages. In their setup, instead of generating pseudo labels the authors generate frame embeddings from a pre-trained ASR model (target-network). Using this technique the authors show improvements over the XLSR model (Conneau et al., 2020a) but the performance is worse compared to the UniSpeech model (Wang et al., 2021). We speculate that XLST performs worse than UniSpeech because UniSpeech optimizes the output representations directly on the ASR task using the CTC loss during the training. For similar reasons XLST performs better than XLSR because XLST use a pre-trained target network whose representations are optimized on the ASR task.

## 4. Discussion

### 4.1. Multilingual vs. monolingual models for low-resource languages

From the papers surveyed, it is clear that multilingual models perform better than monolingual counterparts trained with the same amount of data for a single language. Combining the data of all languages available during pre-training also improves performance compared to using multiple languages only during fine-tuning.

### 4.2. Techniques and architectures

Fine-tuning plays an important role in the accuracy of multilingual models, with multiple studies show-

ing that models fine-tuned to a target language perform better than simply combining data from all languages. Techniques that can be used for further improvements include phone-mapping, using a feature extractor trained on multiple languages and using a common decoder for related languages. Although proximity in terms of language family may play a role in accuracy, some studies show that unrelated languages also benefit from common feature extractors. Multitask Learning setups that use phoneme recognition or Language Identification as auxiliary tasks seem to improve performance only slightly over multilingual models trained jointly with all languages.

### 4.3. Performance on high-resource languages

When building a multilingual model that can be re-used across different languages, it is desirable that the performance on high-resource languages does not degrade while making improvements over low-resource languages. Some studies show that this degradation indeed occurs, however, there are strategies such as sampling and using language ID information that can be used to alleviate the problem. In the case of SSL, it seems that pre-training can help improve performance on high-resource languages as well.

### 4.4. Factors that influence cross-lingual transfer

Many factors affect the overall accuracy on target languages, including the choice of features and the languages that the features are trained on and the choice of architecture while building models trained on multiple languages. Some studies show that performance increases with the number of languages in the model, provided that the model is large and deep enough. In addition to training common feature extractors, it is also seen that training the same decoder for multiple languages performs better if the languages are similar.

### 4.5. Role of SSL and unlabeled data

Although research on this topic is nascent, it is clear that the choice of pre-training data matters. Pre-training from a diverse set of languages, or languages related to the target language is better than restricting it only to one language such as English even if the total pre-training data remains the same. Interestingly, pre-training with a diverse set of languages has been shown to improve performance on languages that are not present in the training data, though more experimentation is needed to study this further. It has also been shown that the pre-training data can be collected from a different domain than the target data and still improve performance. Finally, the size of pre-training and fine-tuning data matters - although it is possible to exploit unlabeled data for pre-training, the fine-tuning data needs to be of a reasonable size to get performance improvements. It is therefore important to invest in a minimum amount of labeled data for fine-tuning for target low-resource languages, unless there is a language

in the multilingual model that is very closely related to the target language.

SSL and the use of unlabeled data are exciting directions for low-resource ASR, particularly when labeled and unlabeled data are from different domains. However, this approach shows some degradation for high-resource languages, that can potentially be addressed by increasing model capacity.

#### 4.6. Open problems and challenges

Although cross-lingual transfer and SSL are promising directions for low-resource language ASR, there are some concerns that still need to be addressed by the research community. More research is needed to determine whether these techniques can be used for very low-resource languages that have only a few hours of data, since labeled data is required for fine-tuning, transfer learning as well as evaluation.

Studies on multilingual ASR have been limited to a few datasets and languages mainly from the Indo-European language family, and more research is needed to determine whether results hold true across languages. Evaluation benchmarks for multilingual ASR should be designed to cover diverse languages in terms of language family and dataset sizes (both labeled and unlabeled). Building such evaluation benchmarks is expensive but would be a crucial investment for improving low-resource language ASR.

Most multilingual models are trained on less than 10 languages and have not reached the size (in terms of languages) of their NLP counterparts that are trained on over 100 languages. An interesting direction of study would be the implications of larger models covering more languages on performance on high resource languages. As models become larger in terms of parameters, training time and training data required, it would be prudent to continue to compare performance with simpler monolingual models, particularly for low-resource languages.

SSL in NLP not only transfers across languages but also across tasks. Transferring representations learned by ASR models trained using SSL is an interesting future direction of research, in which some progress has been made already.

### 5. Conclusion

In this paper, we survey more than 40 papers on multilingual ASR, in which either multiple languages are trained in a single model, or cross-lingual transfer learning is used to improve performance on low-resource languages. Till recently, ASR models were only able to exploit transcribed data, however, with the advent of Self Supervised Learning, they are now also able to use unlabeled data which is an exciting development for low-resource languages. We address questions about whether multilingual models are indeed superior in performance to monolingual models for low-resource languages, as well as for high-resource lan-

guages. We distill key findings from research in multilingual ASR to describe factors that influence cross-lingual transfer and SSL, and provide recommendations for future research.

### 6. References

- Amodei, D., Ananthanarayanan, S., Anubhai, R., Bai, J., Battenberg, E., Case, C., Casper, J., Catanzaro, B., Cheng, Q., Chen, G., et al. (2016). Deep speech 2: End-to-end speech recognition in english and mandarin. In *International conference on machine learning*, pages 173–182. PMLR.
- Baevski, A., Zhou, H., Mohamed, A., and Auli, M. (2020). wav2vec 2.0: A framework for self-supervised learning of speech representations. *arXiv preprint arXiv:2006.11477*.
- Billa, J. (2018). Isi asr system for the low resource speech recognition challenge for indian languages. In *INTERSPEECH*, pages 3207–3211.
- Billa, J. (2021). Improving low-resource asr performance with untranscribed out-of-domain data. *arXiv preprint arXiv:2106.01227*.
- Chan, W., Jaitly, N., Le, Q. V., and Vinyals, O. (2015). Listen, attend and spell. *CoRR*, abs/1508.01211.
- Chung, Y.-A., Zhang, Y., Han, W., Chiu, C.-C., Qin, J., Pang, R., and Wu, Y. (2021). W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training. *arXiv preprint arXiv:2108.06209*.
- Conneau, A., Baevski, A., Collobert, R., Mohamed, A., and Auli, M. (2020a). Unsupervised cross-lingual representation learning for speech recognition. *arXiv preprint arXiv:2006.13979*.
- Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., and Stoyanov, V. (2020b). Unsupervised cross-lingual representation learning at scale. In *ACL*.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22.
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. (2018). BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- Ghoshal, A., Swietojanski, P., and Renals, S. (2013). Multilingual training of deep neural networks. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 7319–7323. IEEE.
- Graves, A., Fernández, S., Gomez, F., and Schmidhuber, J. (2006). Connectionist temporal classification: labelling unsegmented sequence data with

- recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pages 369–376.
- Graves, A., Jaitly, N., and Mohamed, A.-r. (2013). Hybrid speech recognition with deep bidirectional lstm. In *2013 IEEE workshop on automatic speech recognition and understanding*, pages 273–278. IEEE.
- Gulati, A., Qin, J., Chiu, C.-C., Parmar, N., Zhang, Y., Yu, J., Han, W., Wang, S., Zhang, Z., Wu, Y., et al. (2020). Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100*.
- Gupta, A., Chadha, H. S., Shah, P., Chimmwal, N., Dhuriya, A., Gaur, R., and Raghavan, V. (2021). Clsril-23: cross lingual speech representations for indic languages. *arXiv preprint arXiv:2107.07402*.
- Han, W., Zhang, Z., Zhang, Y., Yu, J., Chiu, C.-C., Qin, J., Gulati, A., Pang, R., and Wu, Y. (2020). Contextnet: Improving convolutional neural networks for automatic speech recognition with global context. *arXiv preprint arXiv:2005.03191*.
- Hannun, A., Case, C., Casper, J., Catanzaro, B., Diamos, G., Elsen, E., Prenger, R., Satheesh, S., Sengupta, S., Coates, A., et al. (2014). Deep speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567*.
- Heafield, K. (2011). Kenlm: Faster and smaller language model queries. In *Proceedings of the sixth workshop on statistical machine translation*, pages 187–197.
- Heigold, G., Vanhoucke, V., Senior, A., Nguyen, P., Ranzato, M., Devin, M., and Dean, J. (2013). Multilingual acoustic models using distributed deep neural networks. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 8619–8623. IEEE.
- Hsu, W.-N., Bolte, B., Tsai, Y.-H. H., Lakhotia, K., Salakhutdinov, R., and Mohamed, A. (2021). Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *arXiv preprint arXiv:2106.07447*.
- Huang, J.-T., Li, J., Yu, D., Deng, L., and Gong, Y. (2013). Cross-language knowledge transfer using multilingual deep neural network with shared hidden layers. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 7304–7308. IEEE.
- Javed, T., Doddapaneni, S., Raman, A., Bhogale, K. S., Ramesh, G., Kunchukuttan, A., Kumar, P., and Khapra, M. M. (2021). Towards building asr systems for the next billion users. *arXiv preprint arXiv:2111.03945*.
- Joshi, P., Santy, S., Budhiraja, A., Bali, K., and Choudhury, M. (2020). The state and fate of linguistic diversity and inclusion in the nlp world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293.
- Kahn, J., Lee, A., and Hannun, A. (2020). Self-training for end-to-end speech recognition. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7084–7088. IEEE.
- Kannan, A., Datta, A., Sainath, T. N., Weinstein, E., Ramabhadran, B., Wu, Y., Bapna, A., Chen, Z., and Lee, S. (2019). Large-scale multilingual speech recognition with a streaming end-to-end model. *arXiv preprint arXiv:1909.05330*.
- Khan, S., Pal, M., Basu, J., Samirakshma Bepari, M., and Roy, R. (2018). Assessing Performance of Bengali Speech Recognizers Under Real World Conditions using GMM-HMM and DNN based Methods. In *Proc. 6th Workshop on Spoken Language Technologies for Under-Resourced Languages (SLTU 2018)*, pages 192–196.
- Li, J., Lavrukhin, V., Ginsburg, B., Leary, R., Kuchaiev, O., Cohen, J. M., Nguyen, H., and Gadde, R. T. (2019). Jasper: An end-to-end convolutional neural acoustic model. *arXiv preprint arXiv:1904.03288*.
- Müller, M., Stüker, S., and Waibel, A. (2017). Phonic and graphemic multilingual ctc based speech recognition. *arXiv preprint arXiv:1711.04564*.
- Pratap, V., Sriram, A., Tomasello, P., Hannun, A., Liptchinsky, V., Synnaeve, G., and Collobert, R. (2020). Massively multilingual asr: 50 languages, 1 model, 1 billion parameters. *arXiv preprint arXiv:2007.03001*.
- Sailor, H. B. and Hain, T. (2020). Multilingual speech recognition using language-specific phoneme recognition as auxiliary task for indian languages. In *INTERSPEECH*, pages 4756–4760.
- Shetty, V. M. and NJ, M. S. M. (2020). Improving the performance of transformer based low resource speech recognition for indian languages. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8279–8283. IEEE.
- Srivastava, B. M. L., Sitaram, S., Mehta, R. K., Mohan, K. D., Matani, P., Satpal, S., Bali, K., Srikanth, R., and Nayak, N. (2018). Interspeech 2018 low resource automatic speech recognition challenge for indian languages. In *SLTU*, pages 11–14.
- Synnaeve, G., Xu, Q., Kahn, J., Likhomanenko, T., Grave, E., Pratap, V., Sriram, A., Liptchinsky, V., and Collobert, R. (2019). End-to-end asr: from supervised to semi-supervised learning with modern architectures. *arXiv preprint arXiv:1911.08460*.
- Thomas, S., Ganapathy, S., and Hermansky, H. (2010). Cross-lingual and multi-stream posterior features for low resource lvcsr systems. In *Eleventh Annual Conference of the International Speech Communication Association*.
- Thomas, S., Ganapathy, S., and Hermansky, H. (2012). Multilingual mlp features for low-resource lvcsr systems. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4269–4272. IEEE.



- Tong, S., Garner, P. N., and Boulard, H. (2017). Multilingual training and cross-lingual adaptation on ctc-based acoustic model. *arXiv preprint arXiv:1711.10025*.
- Toshniwal, S., Sainath, T. N., Weiss, R. J., Li, B., Moreno, P., Weinstein, E., and Rao, K. (2018). Multilingual speech recognition with a single end-to-end model. In *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 4904–4908. IEEE.
- Tüske, Z., Pinto, J., Willett, D., and Schlüter, R. (2013a). Investigation on cross-and multilingual mlp features under matched and mismatched acoustical conditions. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 7349–7353. IEEE.
- Tüske, Z., Schlüter, R., and Ney, H. (2013b). Multilingual hierarchical mrasta features for asr. In *Interspeech*, pages 2222–2226.
- Wang, C., Wu, Y., Qian, Y., Kumatani, K., Liu, S., Wei, F., Zeng, M., and Huang, X. (2021). Unispeech: Unified speech representation learning with labeled and unlabeled data. *arXiv preprint arXiv:2101.07597*.
- Watanabe, S., Hori, T., and Hershey, J. R. (2017). Language independent end-to-end architecture for joint language identification and speech recognition. In *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 265–271. IEEE.
- Wu, S. and Dredze, M. (2019). Beto, bentz, becas: The surprising cross-lingual effectiveness of bert. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 833–844.
- Xiong, W., Droppo, J., Huang, X., Seide, F., Seltzer, M., Stolcke, A., Yu, D., and Zweig, G. (2016). Achieving human parity in conversational speech recognition. *arXiv preprint arXiv:1610.05256*.
- Zhang, Z.-Q., Song, Y., Wu, M.-H., Fang, X., and Dai, L.-R. (2021). Xlst: Cross-lingual self-training to learn multilingual representation for low resource speech recognition. *arXiv preprint arXiv:2103.08207*.
- Zhou, S., Zhao, Y., Xu, S., and Xu, B. (2017). Multilingual Recurrent Neural Networks with Residual Learning for Low-Resource Speech Recognition. In *Proc. Interspeech 2017*, pages 704–708.
- Zhou, S., Xu, S., and Xu, B. (2018). Multilingual end-to-end speech recognition with a single transformer on low-resource languages. *arXiv preprint arXiv:1806.05059*.
- Zhou, W., Zeinldeen, M., Zheng, Z., Schlüter, R., and Ney, H. (2021). Acoustic data-driven subword modeling for end-to-end speech recognition. *arXiv preprint arXiv:2104.09106*.