

ACOUSTICS BASED INTENT RECOGNITION USING DISCOVERED PHONETIC UNITS FOR LOW RESOURCE LANGUAGES

Akshat Gupta, Xinjian Li, Sai Krishna Rallabandi, Alan W Black
akshatgu@andrew.cmu.edu, {xinjianl, srallaba, awb}@cs.cmu.edu

Carnegie Mellon University

ABSTRACT

With recent advancements in language technologies, humans are now speaking to devices. Increasing the reach of spoken language technologies requires building systems in local languages. A major bottleneck here are the underlying data-intensive parts that make up such systems, including automatic speech recognition (ASR) systems that require large amounts of labelled data. With the aim of aiding development of spoken dialog systems in low resourced languages, we propose a novel acoustics based intent recognition system that uses discovered phonetic units for intent classification. The system is made up of two blocks - the first block is a universal phone recognition system that generates a transcript of discovered phonetic units for the input audio, and the second block performs intent classification from the generated phonetic transcripts. We propose a CNN+LSTM based architecture and present results for two languages families - Indic languages and Romance languages, for two different intent recognition tasks. We also perform multilingual training of our intent classifier and show improved cross-lingual transfer and zero-shot performance on an unknown language within the same language family.

Index Terms— Intent, Low Resource, Cross Lingual, Multilingual, Long short term Memory

1. INTRODUCTION

In order to bring technology closer to humans, it is important to provide accessible mechanisms of interaction. Speech is regarded as the most natural form of interaction for humans and its accessibility has been aided by improvements in language technologies such as Automatic Speech Recognition and Speech Synthesis. However, lack of annotated data is a major bottleneck for scaling speech technologies to new languages and domains. Therefore, it is useful to design techniques that can perform well in low data scenarios. A fundamental resource required to build such systems is a phonetic lexicon which can translate acoustic input to textual representation. In this paper, we present a novel approach to perform intent recognition purely from acoustics using such a phonetic lexicon. A block diagram representing our system is shown

in Figure 1. Through our approach, we bypass the need to build language specific ASR systems, which are very data intensive, and demonstrate deployable performance for intent recognition using discovered phonetic units.

We test the performance of our system on two language families - Indic and Romance languages, each having a different intent recognition task. We also train our intent classification system multilingually and evaluate its zero-shot performance for a language not in the training set, although still within the same language family. This simulates a zero resource scenario and helps us understand the extent of cross-lingual transfer between languages for our acoustics based intent recognition system. We find that a multilingually trained classification model performs significantly better than a monolingually trained model for an unknown language.

2. RELATED WORKS

Spoken Language Understanding (SLU) systems aim to process spoken utterances for various downstream tasks. Current research in high resourced languages is moving towards building end-to-end SLU systems [1] [2] to eliminate propagating errors through the SLU pipeline. A typical SLU pipeline is made up of two blocks - a speech to text module followed by a natural language understanding (NLU) module. Building speech to text modules require large amounts of labelled speech data, which is scarce for low resourced languages. [3] [4] have previously used outputs of an English ASR system and English phonemes for intent classification in Sinhala and Tamil.

To the best of our knowledge, this paper is the first attempt in literature to use the phonetic units for intent classification. There have been numerous attempts[5, 6] to discover such acoustic units in an unsupervised fashion. In [7], authors presented an approach to modify the speaker diarization system to detect speaker-dependent acoustic units. [8] proposed a GMM-based approach to discover speaker-independent subword units. However, their system requires a separate Spoken Term Detector. Our work is closest to [9] where authors discover symbolic units in an unsupervised fashion for speech to speech translation. Contrary to this work, we employ the

Intents	Number of Utterances
ordering pizza	711
auto-repair appointment	484
order ride service	450
order movie tickets	549
order coffee	292
restaurant reservations	757

Table 1. Class distribution for the Indic dataset.

symbolic units generated by Allosaurus [10] which is trained in a supervised fashion.

3. DATASETS

We study the performance of our acoustics based intent recognition system for two language families - Indic Languages and Romance Languages. For each family we use a different dataset and each language family has a different intent recognition task.

3.1. Dataset for Indic Languages

We use Google’s Taskmaster-1 Dataset [11] for Indic Languages which contains data for user interactions with an autonomous dialogue system collected using the Wizard of Oz methodology [12]. The user dialogues are a written transcripts of the conversations in English. The dataset contains labelled intents and slots for the conversation. We extract the sentence responsible for the labelled intent from the dataset and create an intent recognition dataset. We obtain 3243 utterances in total distributed amongst 6 intents as shown in Table 1. After creating the intent classification dataset, we translate the transcripts in English into four Indic languages - Hindi, Gujarati, Bengali and Marathi, using the Google Translate API. The translated text was used to synthesize audio using the Google Text-To-Speech API for Hindi, Gujarati and Bengali. CLUSTERGEN [13] was used synthesizing for Marathi voice. The voice quality of Marathi is much worse when compared to the other voices generated from Google’s API. The dataset in each language contains two voices - one male and one female. The audios are then passes into Allosaurus [10] to discover phonetic units and create a phonetic transcription of the audio.

3.2. Dataset for Romance Languages

To work with Romance languages, we create an intent recognition dataset from the MultiWoz dataset [14]. The dataset contains a large number of dialogues between humans and robots where each utterance is associated with a json object containing the conversational context. The context has rich information about the intent of humans. The largest context

Intents	Number of Utterances
Monday	743
Tuesday	718
Wednesday	757
Thursday	738
Friday	763
Saturday	779
Sunday	806

Table 2. Class distribution for the Romance dataset.

class in the dataset is about the reservation day whose distributions are shown in the Table 2. This class is used as our Romance languages dataset. This dataset is then prepared in a similar way as done for the Indic language, where we translate the original English utterances into 4 different Romance languages - Italian, Portuguese, Romanian and Spanish. The translated text is synthesized with the Google TTS engine, and then transcribed into phonetic units with Allosaurus [10].

4. MODELS

A block diagram depicting our acoustics based intent recognition system utilizing a phonetic transcription is shown in Figure 1. The input audio is directly fed into a system that can generate hypothesized phonetic units. For our work, we use the Allosaurus library [10] which is a nearly-universal phone recognition system. For this work, we employ the language dependent phones, which basically means we’re providing an identifier to Allosaurus for audio language. The phonetic transcription is then sent to an intent classifier that does the classification purely based on the generated sequence of phones. Such very simple systems can be used to build powerful tools, especially for low resource languages, as shown in [15].



Fig. 1. Block Diagram showing a general acoustics based intent recognition system.

We use a Naive Bayes classifier as our baseline with add-1 smoothing and absolute discounting. We also propose a neural network architecture shown in Figure 2 to compare with the baseline results. The architecture is based on LSTMs (long-short term memory) [16] for modeling sequential information where the contextual information is encoded using CNNs (convolutional neural network).

The input to the network is a sequence of phones $\mathbf{x} = x_1, x_2, \dots, x_t$, where each phonetic unit is passed to a 128 dimensional embedding vector. The embedding layer converts the input sequence into a dense vector representation which is

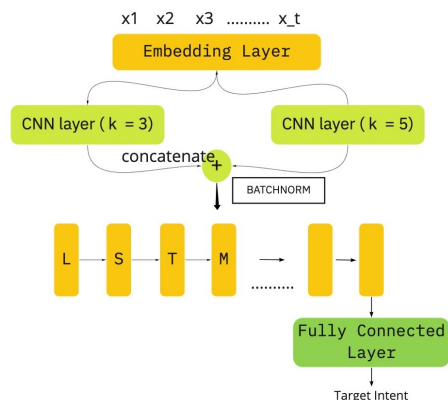


Fig. 2. Block Diagram depicting the architecture of our proposed neural network.

then sent to two 1-d CNN layers of kernel size $k = 3, 5$. The CNN layers have 128 filters and capture trigram and 5-gram features from the phonetic transcription. The outputs of each of the CNN layers are concatenated to create a 256 dimension long embedding vector where each embedding vector now has contextual information encoded in it.

The concatenated embeddings are passed through the LSTM layer consisting of 128 neurons. The hidden state of the LSTM layer at the final time step is sent to a linear layer for intent classification.

5. EXPERIMENTS

We test our acoustics based intent recognition system for two sets of languages across two different language families - Indic and Romance language families. We perform monolingual and multilingual training for both baseline and our proposed neural network architecture and test the model performance for multiple languages.

5.1. Monolingual Training Results

In this section we present results for intent classification architectures trained on a single language. Table 3 presents the classification results for Indic languages and Table 5 for Romance languages. The diagonal elements in the tables show the classification accuracy for training and testing performed on the same language. The numbers in the bracket show performance with the baseline (Naive Bayes) classifier. We see that our proposed neural network architecture improves on our baseline significantly.

Cross-lingual testing results for monolingually trained classification models are also shown in Tables 3 and 5. The performance is relatively poor when the classification model is trained on only one language due to minimal cross-lingual transfer. Language pairs for linguistically similar languages show higher performance. This can be seen for language

Test Train	Hin	Guj	Mar	Ben
Hin	92.0 (89.3)	54.7(59.7)	43.7(36.7)	54.3(45.3)
Guj	52.3(50.3)	93.3 (91.7)	52.0(47.0)	63.0(39.3)
Mar	52.0(35.0)	66.3(49.7)	87.7 (84.3)	58.0(37.0)
Ben	48.0(41.7)	54.7(38.3)	45.7(31.3)	95.0 (93.0)

Table 3. Classification Accuracy for monolingual training for Indic Languages - Hindi (Hin), Gujarati (Guj), Marathi (Mar) and Bengali (Ben). The numbers in the bracket are the baseline results using a Naive Bayes classifier.

Test Train	Hin	Guj	Mar	Ben
HGM	85.3(84.7)	90.3(86)	75.6(78.3)	80.7 (58.3)
HGB	87.3(84)	90.0(84)	61.7 (54.3)	90.3(89.3)
HMB	84.3(88.7)	62 (65.4)	80.7(76.6)	88.3(88)
GMB	65.3 (63.7)	86.7(84.7)	83.0(80)	92.0(89.7)

Table 4. Average Classification Accuracy for a multilingually trained model. The languages in bold are the languages that are not present in the train set. The numbers in the bracket are the baseline results using a Naive Bayes classifier.

pairs Hindi-Gujarati and Gujarati-Marathi in Indic language family and pairs Italian-Portuguese and Italian-Spanish in the Romance language family. These language pairs are also geographically close. The cross lingual results are in general better for the Indic Dataset when compared to the Romance dataset. We believe this is because all Indic languages have some amount of code mixing within them. Therefore, there is a larger cross-lingual transfer of features between any pair of languages in the Indic language family when compared to the Romance language.

5.2. Multilingual Training Results

With the aim of improving performance on a language not present in our training set and simulating a zero resource scenario, we train a multilingual model. The training set size is kept the same and the exact same train-test split is used for accuracy scores as used for monolingual results. Let $T = [L_1, L_2, \dots, L_n]$ be the set of languages we use to train the classifier. Then the training set is divided randomly and equally amongst the ' n ' languages present in the training set.

The results for multilingual training can be seen in Table 4 and Table 6 when trained on $n = 3$ languages. The results in bold are for the language not in the training set. The numbers in the bracket show performance with the baseline (Naive Bayes) classifier. We see that our proposed neural network architecture improves on the baseline results significantly in almost all cases. The power of multilingual training becomes apparent when we look at the performance on a language not present in T. We find that a multilingual classi-

Test Train	Ita	Por	Ron	Spa
Ita	88.6 (82.4)	27.6(24.2)	28.6(33.9)	39.0(32.5)
Por	30.6(23.3)	88.6 (74.6)	28.2(22.2)	36.0(26.2)
Ron	32.6(31.6)	29.6(17.4)	86.6 (76.5)	46.0(33.3)
Spa	46.1(35.8)	43.9(40.2)	35.4(33.9)	88.7 (83.3)

Table 5. Classification Accuracy for monolingual training for Romance Languages - Italian (Ita), Portuguese (Por), Romanian (Ron and Spanish (Spa). The numbers in the bracket are the baseline results using a Naive Bayes classifier.

Test Train	Ita(I)	Por(P)	Ron(R)	Spa(S)
IPR	87.0(78.8)	88.9(62.1)	85.4(69.1)	60.4 (43.9)
IPS	88.6(77.2)	88.9(62.9)	37.4 (40.7)	88.9(80.4)
IRS	88.1(88.2)	41.3 (30.0)	86.6(73.2)	88.7(80.0)
PRS	50.3 (40.9)	87.8(59.9)	84.9(69.1)	88.3(79.9)

Table 6. Classification Accuracy for a multilingually trained model. The languages in bold are the languages that are not present in the train set. The numbers in the bracket are the baseline results using a Naive Bayes classifier.

fier always performs better on an unknown language $L_u \notin T$ when compared to zero-shot transfer by monolingual model without significant performance loss in individual languages. Its important to consider that we haven't augmented the data in any form, thus the multilingual model see far fewer example of a specific language than monolingual models.

The results in Table 4 and Table 6 show that there is larger amount of cross-lingual transfer when the model is trained on many languages from the same language family. In a practical scenario, this means that a deployed multilingual model is more likely to generalize better to an unknown language or variation in dialect than a monolingual model. This is especially useful for the case of low-resource languages for which it's hard to collect any training data. We illustrate this point by taking the example of Bengali in the Indic language family. There are four models across Tables 3 and 4 that do not have Bengali in the training set. These are the models not highlighted in the Bengali column in Table 3 and the highlight value in the Bengali column in Table 4. The multilingual model performs best out of these four models. This is also true for Spanish in Tables 5 and 6.

The performance for an unknown language $L_u \notin T$ can further be improved by injecting a very small amount of data for L_u in the training set. We added training data for language L_u in increments of a ratio of 0.05 of the training set as shown in Figure 3. We see that introducing even the slightest amount of training data for the unknown language increases its performance significantly while not affecting the performance of the other languages. Figure 3 shows an increase in performance of about 9% for Marathi, 14% for Hindi and

17% for Gujarati only by an injection of data 5% the size of training dataset.




Fig. 3. Plot showing performance of a multilingual intent classification model for when data for a language is injected into the training set in increments of ratio of 0.05 for Indic languages. For example, HGM -> B represents a model trained on Hindi, Gujarati, Marathi and we're checking the increase in performance on Bengali by injecting Bengali data into the training dataset.

6. DISCUSSION

We present a novel approach for intent recognition in low resourced languages with experiments on two different language families. It was shown that zero-shot performance for a language not in the training set of the model but still within the language family can be improved with multilingual training. This helps in maximal cross-lingual transfer between languages that are linguistically and geographically closer to each other. We also found that performance for such a language not in the training set can be improved significantly by introducing a minimal amount of training data.

Our present work was based on synthesized data due to the absence of enough natural speech datasets for intent recognition for low resource languages. Future work can include corroboration of our results with natural speech. The synthesized speech also had little speaker variation in terms of speaker style or prosody though we did include variation in speaker gender.

7. CONCLUSION

We present a novel acoustics based intent recognition system that classifies intents from phonetic transcripts generated using a (nearly-)universal phone recognizer, bypassing the need to build language specific ASR. We also show that multilingual training within same language families produce better zero shot transfer within same the language family when compared to monolingual models.

8. REFERENCES

- [1] Loren Lugosch, Mirco Ravanelli, Patrick Ignoto, Vikrant Singh Tomar, and Yoshua Bengio, “Speech model pre-training for end-to-end spoken language understanding,” *arXiv preprint arXiv:1904.03670*, 2019.
- [2] Martin Radfar, Athanasios Mouchtaris, and Siegfried Kunzmann, “End-to-end neural transformer based spoken language understanding,” *arXiv preprint arXiv:2008.10984*, 2020.
- [3] Yohan Karunanayake, Uthayasanker Thayasivam, and Surangika Ranathunga, “Transfer learning based free-form speech command classification for low-resource languages,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, 2019, pp. 288–294.
- [4] Yohan Karunanayake, Uthayasanker Thayasivam, and Surangika Ranathunga, “Sinhala and tamil speech intent identification from english phoneme based asr,” in *2019 International Conference on Asian Language Processing (IALP)*. IEEE, 2019, pp. 234–239.
- [5] Leonardo Badino, Claudia Canevari, Luciano Fadiga, and Giorgio Metta, “An auto-encoder based approach to unsupervised learning of subword units,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 7634–7638.
- [6] Janek Ebbens, Jahn Heymann, Lukas Drude, Thomas Glarner, Reinhold Haeb-Umbach, and Bhiksha Raj, “Hidden markov model variational autoencoder for acoustic unit discovery,” in *INTERSPEECH*, 2017, pp. 488–492.
- [7] Marijn Huijbregts, Mitchell McLaren, and David Van Leeuwen, “Unsupervised acoustic sub-word unit detection for query-by-example spoken term detection,” in *2011 IEEE international conference on Acoustics, speech and signal processing (ICASSP)*. IEEE, 2011, pp. 4436–4439.
- [8] Aren Jansen, Samuel Thomas, and Hynek Hermansky, “Weak top-down constraints for unsupervised acoustic model training,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2013, pp. 8091–8095.
- [9] Steven Hillis, Anushree Prasanna Kumar, and Alan W Black, “Unsupervised phonetic and word level discovery for speech to speech translation for unwritten languages,” in *INTERSPEECH*, 2019, pp. 1138–1142.
- [10] Xinjian Li, Siddharth Dalmia, Juncheng Li, Matthew Lee, Patrick Littell, Jiali Yao, Antonios Anastasopoulos, David R Mortensen, Graham Neubig, Alan W Black, et al., “Universal phone recognition with a multilingual allophone system,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 8249–8253.
- [11] Bill Byrne, Karthik Krishnamoorthi, Chinnadhurai Sankar, Arvind Neelakantan, Daniel Duckworth, Semih Yavuz, Ben Goodrich, Amit Dubey, Andy Cedilnik, and Kyu-Young Kim, “Taskmaster-1: Toward a realistic and diverse dialog dataset,” *arXiv preprint arXiv:1909.05358*, 2019.
- [12] Bruce Hanington and Bella Martin, *Universal methods of design: 100 ways to research complex problems, develop innovative ideas, and design effective solutions*. Rockport Publishers, 2012.
- [13] Alan W Black and Prasanna Kumar Muthukumar, “Random forests for statistical speech synthesis,” in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [14] Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Inigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić, “Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling,” *arXiv preprint arXiv:1810.00278*, 2018.
- [15] Akshat Gupta, Sai Krishna Rallabandi, and Alan W Black, “Mere account mein kitna balance hai?—on building voice enabled banking services for multilingual communities,” *arXiv preprint arXiv:2010.16411*, 2020.
- [16] Sepp Hochreiter and Jürgen Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.