



Allophant: Cross-lingual Phoneme Recognition with Articulatory Attributes

Kevin Glocker¹, Aaricia Herygers¹, Munir Georges^{1,2}

¹Almotion Bavaria, Technische Hochschule Ingolstadt, Germany

²Intel Labs, Germany

firstname.lastname@thi.de

Abstract

This paper proposes Allophant, a multilingual phoneme recognizer. It requires only a phoneme inventory for cross-lingual transfer to a target language, allowing for low-resource recognition. The architecture combines a compositional phone embedding approach with individually supervised phonetic attribute classifiers in a multi-task architecture. We also introduce Allophoible, an extension of the PHOIBLE database. When combined with a distance based mapping approach for grapheme-to-phoneme outputs, it allows us to train on PHOIBLE inventories directly.

By training and evaluating on 34 languages, we found that the addition of multi-task learning improves the model’s capability of being applied to unseen phonemes and phoneme inventories. On supervised languages we achieve phoneme error rate improvements of 11 percentage points (pp.) compared to a baseline without multi-task learning. Evaluation of zero-shot transfer on 84 languages yielded a decrease in PER of 2.63 pp. over the baseline.

Index Terms: speech recognition, cross-lingual, zero-shot, phoneme recognition

1. Introduction

Speech technologies such as automatic speech recognition (ASR) have greatly improved in recent years. However, despite the increase in research interest [1, 2, 3, 4, 5, 6], for many low-resource languages these technologies are still not available or perform poorly. Such languages may be endangered or lack a “stable orthography” [7]. Similarly, many models experience difficulties in recognizing regional or non-native accented speech [8, 9, 10], for which there is often little data available.

To reduce the amount of training data needed for these languages, various techniques (e.g., probabilistic transcriptions [11]) and benchmarks [12] have been proposed. To avoid the dependence on training data in the target language, architectures have been introduced that allow zero-shot phoneme recognition on unseen languages using only their phoneme inventories [13, 14]. Both architectures use allophone-to-phoneme mappings (i.e., variations of phonemes that retain word meanings but are generally articulated in distinct contexts [15]).

In the ASR2K architecture [16], multilingual acoustic and pronunciation models are trained jointly without any supervision. They successfully use pre-trained self-supervised learning (SSL) models, such as the multilingually pre-trained XLS-R [17, 16]. Additionally, multi-task learning with articulatory attributes has been shown to improve character error rates in Mandarin [18] and cross-lingual phoneme error rates [19].

This work was partially supported by the German Federal Ministry of Education and Research (BMBF) under funding nr. 16DHBKI032

Following the examples set by Allosaurus [13], a “multilingual allophone system”, and AlloVera [20], a “multilingual allophone database”, we thus present our multilingual phoneme recognition architecture called “Allophant” and phoneme database called “Allophoible”. Allophant combines an allophone layer and phonetic composition [21, 16] with multi-task learning for cross-lingual phoneme recognition.

We describe the Allophant architecture in Section 2. In Section 3 we present Allophoible, an extension of the PHOIBLE database [22] that includes articulatory attributes for additional phoneme segments such as allophones. We describe our experiments in Section 4 and present and discuss the results in Section 5 before concluding in Section 6.

2. Allophant Architecture

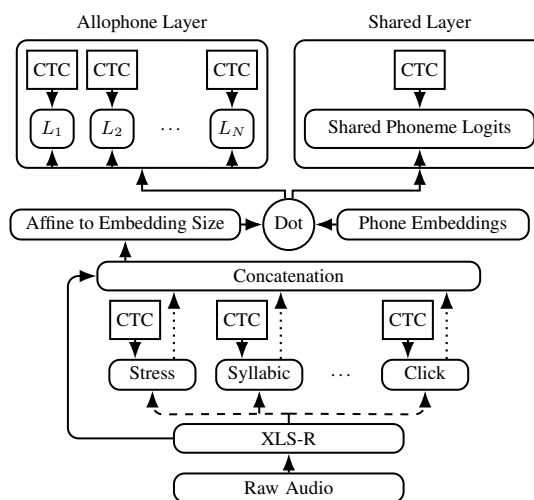


Figure 1: Illustration of the Allophant architecture, including an allophone layer [13], phone composition [14] and hierarchical multi-task connections [19]. Multi-task models use the dashed connections and hierarchical multi-task models also use the dotted connections. “Shared” models use the shared layer while all other models use the allophone layer instead.

The proposed Allophant architecture and its evaluated variants are illustrated in Figure 1. As in previous work on zero-shot phoneme recognition [16, 23], we finetune a model trained with an SSL objective. We choose the 300 million parameter version of the cross-lingual speech representation model (XLS-R), which was pre-trained on 128 languages [23]. XLS-R was found to outperform other SSL models for this task [16].

As the phoneme classifier, we use the compositional architecture from [14] and apply it to the PHOIBLE feature set [22]

as motivated in Section 3. In this architecture, phone embeddings are computed by summing the embeddings of their attribute values. We designate an embedding to represent each possible value of the 35 articulatory attributes in addition to a special “blank” attribute, which represents blanks for connectionist temporal classification (CTC) training [24, 14].

Each attribute can have a value of “+” if it exists, “-” if it does not, or “0” if it is impossible in combination with other attributes of a phoneme. The only exception for this is the stress attribute, which is only “-” or “0” in PHOIBLE. Since we do not consider tonal languages (e.g., Mandarin) in this paper, the tone attribute was unused and therefore removed from training and inference. Complex phoneme segments such as some diphthongs, affricates, and multiple articulations contain attribute contours such as “+,-,-” instead of only a single value. For simplicity, we only use the first attribute from complex contours for the composition.

Out of 107 attribute values, six are not used by any phonemes in the training data and are replaced by zero vectors to not affect inference. The missing attribute values are [+retractedTongueRoot], [+advancedTongueRoot], [+epilaryngealSource], [+raisedLarynxEjective], [+loweredLarynxEjective], and [+click]. We then compute phone logits with the scaled dot product [25, 14] between the output hidden representation from the final XLS-R layer and the composed phone or blank embeddings. At training time, phoneme probabilities are computed by additionally passing the phone logits through an allophone layer [13, 14, 16]. The resulting language-specific phoneme sequences are then trained with CTC loss.

Following our previous work [19], we additionally explicitly supervise articulatory attribute classifiers for all 35 attributes that are also used for embedding composition. In this architecture, CTC loss is computed for each attribute sequence individually and minimized together with the phoneme CTC loss during training. In contrast to embedding composition, we use the full complex contours in label sequences for attribute classifiers. As a result, attribute sequences may differ in length.

To analyze the interactions between embedding composition, multi-task learning and the allophone layer, we test five variations of the architecture. All variants use phoneme composition. Two different architectures that are trained only on phoneme recognition without multi-task learning are used as baselines. Our first baseline is closest to the compositional model first introduced in [14] and used as the acoustic model in [16]. We also consider a simplified form of the proposed architecture without an allophone layer to better understand the interaction of multi-task learning with phoneme composition on its own. In variants using this “Shared” layer instead we compose embeddings for the union of all phonemes that occur in the training languages instead of all allophones during training and optimize it directly. This configuration is comparable to shared phoneme models used in previous work [13], where a feed-forward projection layer was used instead of embedding composition to compute phoneme logits [23, 19].

In addition, we evaluate two more multi-task architectures to identify how the attribute classifiers impact phoneme recognition. “Multi-Task Shared” is the multi-task analogue to “Baseline Shared” for direct comparison without the allophone layer. “Multi-Task Hierarchy” adds the hierarchical connection between attribute classifiers and the phoneme classifier as in [19]. In this architecture, probability distributions from attribute classifiers are concatenated with the acoustic model outputs from XLS-R before passing the resulting vector through the affine transformation to embedding size.

Table 1: Zero-shot transfer results on UCLA including average PERs, AERs and the variances of PERs between languages.

Name	PER	PER σ^2	AER
Baseline (15k steps)	57.01%	374.35	–
Baseline Shared	48.25%	250.59	–
Multi-Task Shared	46.05%	250.34	19.52%
Multi-Task	45.62%	229.06	19.44%
Multi-Task Hierarchy	46.09%	230.28	19.18%

3. Allophoible: Allophone Inventory

Allophoible¹ (version 1.0) is an open-access inventory that builds upon PHOIBLE 2.0 [22], which contains phoneme information for ≈ 2200 languages. We provided phoneme attributes for diacritics and phonemes that were not included as phonemes in PHOIBLE 2.0 but were included as allophones or in eSpeak NG [26] (version 1.51). This allows us to use the compositional architecture from [14] with PHOIBLE allophone mappings and compose phoneme embeddings for eSpeak NG phonemes.




Figure 2: Absolute changes in PER on UCLA (top) and CV (bottom) by language family (Glottolog [27]) from “Baseline Shared” to “Multi-Task” learning in percentage points.

4. Experiments

4.1. Datasets

We selected a subset of at most 800,000 utterances per language from 34 eSpeak-supported, non-tonal languages from the Mozilla Common Voice (CV) corpus [28] (version 10.0). In total, our training set consisted of 4628h of speech. For validation and testing we used the CV development and test sets and filtered some utterances as described in Section 4.2.

Cross-lingual transfer was evaluated on the phonetically transcribed utterances from the first release of the UCLA Phonetic Corpus [29]. It contains 5,509 utterances from 95 low-

¹<https://github.com/Aariciah/allophoible/releases/tag/v1.0.0>

resource languages, of which 11 occurred in our training set. Since evaluation on these languages would not be zero-shot, they were removed. Their ISO 639-3 codes are: “ben”, “ces”, “dan”, “ell”, “eus”, “fin”, “hin”, “hun”, “lit”, “mlt”, and “nld”.

Furthermore, from the remaining 84 languages, 17 were included in the pre-training data for XLS-R. While the model being aware of the acoustic properties of these languages could benefit phoneme recognition after fine-tuning, we did not find any significant effect (Welch’s T-test: $t(22.7) = 0.64, p = 0.528$ for “Multi-Task”). Their ISO 639-3 language codes are: “abk”, “afr”, “asm”, “guj”, “hau”, “haw”, “heb”, “hrv”, “hye”, “isl”, “kan”, “khm”, “lav”, “lug”, “mal”, “mya”, “yue”. For testing, we use the phoneme inventories specified in the corpus to compose phoneme embeddings for each language.

4.2. Data and Processing

We used the allophone inventories available in PHOIBLE [22] to provide allophone mappings for the architectures using an allophone layer. This is applied to all languages in the training set. Since the phoneme inventories in the database differ to varying degrees from the grapheme-to-phoneme output, we mapped phonemes to their closest equivalents in each inventory.

All phonemes that are already in the target inventory were retained. Similarly to previous work [23], we then mapped the remaining phonemes to the target phonemes with the lowest attribute Hamming distance. For simplicity, we only considered the first attribute in contours of complex segments for this mapping. Finally, we prevented e.g., diphthongs being mapped to single vowels by splitting complex segments whenever the number of segments between pairs of matched phonemes differ. After splitting, each sub-segment was mapped individually. Using this approach, we covered approximately 79% of PHOIBLE inventories for each language compared to 17.5% without it.

To provide a phoneme transcription of the CV data we used eSpeak NG, which was shown to outperform other grapheme-to-phoneme models [23]. Characters in eSpeak NG that were not International Phonetic Alphabet characters were manually replaced. For eSpeak NG, we constructed inventories from the union of all phonemes that occur in the transcriptions of each language. In all three CV subsets (i.e., train, validation, and test) we removed the segments with mixed scripts (e.g., Latin character text in Tamil sentences) as eSpeak NG treats them as being in a different language. This would lead to a mix in phoneme inventories (e.g., English phonemes in the Tamil inventory). To ensure that evaluation results are not skewed based on our mapping, we test on the unmapped reference transcriptions. For this, we computed embeddings for all phonemes that occur in the training and test set for all languages at test time.

4.3. Training

Allophant was implemented in Python with PyTorch [30] and TorchAudio [31]. Our implementation is open-source². We use the pre-trained XLS-R model provided by Hugging Face with the transformers library [32, 17]. We use a dropout rate of 0.2 after the final XLS-R output layer. The attribute embedding size for composition is set to 640 as in previous work [14]. We use dynamic batching where batches are constructed from as many raw audio frame sequences as can fit in an at most 16,000,000 element matrix, including padding.

Following [17], we use language upsampling with $\alpha = 0.5$ for more diverse multilingual training batches. We use

²<https://github.com/kgnlp/allophant>

Table 2: CV test set results for languages in the training data.

Name	PER	PER σ^2	AER
Baseline (15k steps)	46.95%	133.31	–
Baseline Shared	45.35%	140.17	–
Multi-Task Shared	41.20%	153.31	8.88%
Multi-Task	34.34%	149.25	8.36%
Multi-Task Hierarchy	34.35%	128.86	8.56%

transformer-style warmup [25] but we keep the learning rate constant for 10,000 steps after linearly increasing it for 2,500 before starting the decay. This change and step numbers are inspired by the schedule from [23]. Also following their work, we keep the feature encoder of XLS-R frozen. However, we finetune all other layers for the entire training process. We finetuned our models for 30,000 updates on single NVIDIA A100 instances with 40gb of memory. The average training duration was ≈ 60 hours. Since not all utterances are seen after 30,000 updates, we ensured that batches were sampled with the same random seed.

5. Results & Discussion

Results for zero-shot transfer are presented in Section 5.1 and supervised languages analyzed in Section 5.2. We use average attribute error rate (AER) for evaluation. It is computed by calculating the average of individual attribute error rates as in [19]. The “Baseline” model was evaluated after 15,000 (15k) steps instead of 30,000 to allow for a fairer comparison since it started to overfit to the training data. PERs after 30,000 steps increased by 7.5 pp. and 3.3 pp. on UCLA and CV respectively.

5.1. Zero-Shot Transfer on the UCLA Phonetic Corpus

Evaluation results of the zero-shot capabilities of our model on the UCLA Phonetic Corpus are listed in Table 1. “Baseline Shared” performs substantially better than “Baseline” by 8.76 percentage points (pp.) PER. This suggests that the combination of composition and allophone layer does not generalize as well.

An explanation for this potential overfitting is that the allophone layer allows the model to maximize logits of allophones of a phoneme even if its attributes do not represent the acoustic signal well enough. The small difference between “Baseline” and “Baseline Shared” on the training data presented in Section 5.2 supports this hypothesis. Potentially, the initialization of the attribute embeddings lead to some allophones that have a high feature Hamming distance to the corresponding phoneme to start with high logits. Since only the most probable allophone is selected in the allophone layer via max pooling [13], overfitting to initially high scoring allophones is possible.

Both non-hierarchical “Multi-Task” models outperform our strongest baseline model (“Baseline Shared”). The “Multi-Task Shared” model lowers the PER by 2.2 pp. and “Multi-Task” by 2.63 pp., showing the overall benefit of explicitly supervising attribute classifiers to improve the generalization of unseen languages. Compared to “Baseline”, adding multi-task learning decreases the PER substantially by 11.39 pp. without the overfitting issues. We suspect that attribute level supervision signals indirectly help stabilize the attribute embeddings in the composition layer. This might prevent them from diverging too much from underlying attributes to maximize specific allophones.

We did not encounter the same benefits of using hierarchi-




Figure 3: “Baseline Shared” and “Multi-Task” PERs for CV test languages with corresponding hours of training data on a log scale.

cal multi-task learning as in our previous work [19]. With a PER difference of 0.47 pp, it performs almost identically. This might be since in contrast to a projection layer with independent weights for each phoneme, composed phone embeddings do not suffer from the same sparsity and improvements of attribute embeddings for one phoneme also benefits others. This sharing of attribute-level information could only be achieved by the hierarchical connection in [19]. On the attribute level, we found that AERs are within ≈ 0.3 of each other across architecture variations. While the hierarchical models reach the lowest AERs, the difference is smaller than in our previous work [19] likely due to the same effects we observe on a phoneme level.

In contrast to [19], where a feed-forward layer was used instead of embedding composition and no correlation between PER and AER was found ($r^2 = 0.016$), we find a moderate correlation for “Multi-Task” ($r^2 = 0.679$). This shows, how acoustic representations that benefit attribute classification are more likely to also directly improve phoneme recognition with the compositional approach.

5.2. Common Voice

Results on the CV tests sets for languages in the training data are shown in Table 2. We find overall very similar patterns in how models perform relative to each other as on UCLA. When comparing the baselines, the difference between the “Baseline” and “Baseline Shared” models of 1.6 PER is much smaller than on UCLA. This provides further evidence for its tendency to overfit to the training data and languages.

As in the zero-shot results, both multi-task models outperform the baselines. Compared to the baselines, the impact of multi-task learning is larger on CV than on unseen languages. “Multi-Task Shared” outperforms “Baseline Shared” by 4.15 pp. PER. Our “Multi-Task” model lowers PER substantially by 11 pp. over “Baseline Shared” and is our best model overall. This shows that multi-task learning benefits the embedding composition layers through improved acoustic representations both cross-lingually and on supervised languages.

We find that “Multi-Task Hierarchical” yields almost identical PERs as “Multi-Task”, with only a decrease in PER variance across languages. This shows that the hierarchical connection does not benefit supervised languages either, likely for the same reasons outlined in Section 5.1. For higher resource languages, additional resources such as phoneme n-gram language models [23] have been shown to decrease PERs further. Investigating this is left for future work. Correlation between PER and AER is lower on CV than UCLA for “Multi-Task” ($r^2 = 0.571$).

A more detailed comparison of error rates across languages between “Baseline Shared” and “Multi-Task” model alongside

hours of training data can be seen in Figure 3. “Multi-Task” reduces PERs to varying degrees for every language in CV. We improve recognition the most on Kyrgyz (ky, 38.8 pp.), English (en, 27.4 pp.), French (fr, 24.9 pp.), and Polish (pl, 21.9 pp.). To further analyze the difference in recognition performance, the relative changes are shown by language family in Figure 2. Recognition of the three Turkic languages improved the most on average, mainly through Kyrgyz and Turkish. It is followed by the 22 Indo-European languages, Basque, and Tamil. The latter is the single Dravidian language in our training set.

In UCLA, Atlantic-Congo is the language family with the most improvements, likely due to the PER improvements of 10 pp. in Swahili (sw). Afro-Asiatic languages likely also benefit from the 3.6 pp. PER improvement of Maltese (mt). Furthermore, PER improvements for Austroasiatic and Austronesian languages can be explained by the improved modeling of the single languages from these families in our CV training set. In some language families, PER also increases after Multi-Task learning, which requires further investigation in future work. In particular, for the single Salishan language in UCLA, PER rises by almost 10 pp. over “Baseline Shared”.

6. Conclusion

We proposed Allophant, a multilingual phoneme recognizer that can be applied zero-shot to low-resource languages. Its architecture combines phonetic embedding composition with multi-task articulatory attribute classifiers, improving phoneme recognition across both unseen and supervised languages. We also introduced Allophoible, an extension of the PHOIBLE database with attributes for all of its allophones. Together with a mapping scheme of phonemes from eSpeak NG to PHOIBLE inventories, this allowed us to train models with a previously introduced attribute composition and allophone layer on 34 languages.

Additionally, our multi-task learning model outperformed our strongest baseline without it by 2.63 pp. PER on cross-lingual transfer to 84 languages of the UCLA phonetic corpus. Furthermore, our multi-task model outperformed the baseline by 11 pp. PER on the supervised languages from the CV corpus. We showed that multi-task learning compliments the attribute embedding composition architectures by improving their generalization to unseen inventories and reducing overfitting.

A limitation of this work is that phoneme inventories with attributes must be available for the target languages. Future work could expand on inferring inventories for unseen languages [33]. More work is needed to investigate the effects of the Allophant architecture on the recognition of tonal languages and regional or non-native language variants.

7. References

- [1] O. Scharenborg, F. Ciannella, S. Palaskar, A. Black, F. Metze, L. Ondel, and M. Hasegawa-Johnson, "Building an ASR system for a low-resource language through the adaptation of a high-resource language ASR system: Preliminary results," in *Proc. IC-NLSSP*, 2017.
- [2] A. Magueresse, V. Carles, and E. Heetderks, "Low-resource languages: A review of past work and future challenges," *arXiv preprint arXiv:2006.07264*, 2020.
- [3] H. Kamper, Y. Matushevych, and S. Goldwater, "Multilingual acoustic word embedding models for processing zero-resource languages," in *ICASSP 2020*, 2020, pp. 6414–6418.
- [4] C. Yu, M. Kang, Y. Chen, J. Wu, and X. Zhao, "Acoustic modeling based on deep learning for low-resource speech recognition: An overview," *IEEE Access*, vol. 8, pp. 163 829–163 843, 2020.
- [5] H. Yadav and S. Sitaram, "A survey of multilingual models for automatic speech recognition," in *Proc. 13th LREC*. Marseille, France: European Language Resources Association, Jun. 2022, pp. 5071–5079.
- [6] E. van der Westhuizen, H. Kamper, R. Menon, J. Quinn, and T. Niesler, "Feature learning for efficient ASR-free keyword spotting in low-resource languages," *Computer Speech & Language*, vol. 71, p. 101275, 2022.
- [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] P.-S. Huang and M. Hasegawa-Johnson, "Cross-dialectal data transferring for Gaussian mixture model training in Arabic speech recognition," in *4th International Conference on Arabic Language Processing*, 2012.
- [9] S. Feng, O. Kudina, B. M. Halpern, and O. Scharenborg, "Quantifying bias in automatic speech recognition," *arXiv preprint arXiv:2103.15122*, 2021.
- [10] A. Herygers, V. Verkhodanova, M. Coler, O. Scharenborg, and M. Georges, "Bias in Flemish automatic speech recognition," in *Studientexte zur Sprachkommunikation: Elektronische Sprachsignalverarbeitung 2023*, C. Draxler, Ed. TUDpress, Dresden, 2023, pp. 158–165.
- [11] M. A. Hasegawa-Johnson, P. Jyothi, D. McCloy, M. Mirbagheri, G. M. d. Liberto, A. Das, B. Ekin, C. Liu, V. Manohar, H. Tang, E. C. Lalor, N. F. Chen, P. Hager, T. Kekona, R. Sloan, and A. K. C. Lee, "Asr for under-resourced languages from probabilistic transcription," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 1, pp. 50–63, 2017.
- [12] T. A. Nguyen, M. de Seyssel, P. Rozé, M. Rivière, E. Kharitonov, A. Baevski, E. Dunbar, and E. Dupoux, "The Zero Resource Speech Benchmark 2021: Metrics and baselines for unsupervised spoken language modeling," in *NeurIPS Workshop on Self-Supervised Learning for Speech and Audio Processing*, Dec. 2020.
- [13] X. Li, S. Dalmia, J. Li, M. Lee, P. Littell, J. Yao, A. Anastasopoulos, D. R. Mortensen, G. Neubig, A. W. Black, and F. Metze, "Universal phone recognition with a multilingual allophone system," in *ICASSP 2020*, 2020, pp. 8249–8253.
- [14] X. Li, J. Li, F. Metze, and A. W. Black, "Hierarchical Phone Recognition with Compositional Phonetics," in *Proc. Interspeech 2021*, 2021, pp. 2461–2465.
- [15] P. Ladefoged and K. Johnson, *A Course in Phonetics*. Cengage Learning, 2014.
- [16] X. Li, F. Metze, D. R. Mortensen, A. W. Black, and S. Watanabe, "ASR2K: Speech Recognition for Around 2000 Languages without Audio," in *Proc. Interspeech 2022*, 2022, pp. 4885–4889.
- [17] A. Babu, C. Wang, A. Tjandra, K. Lakhota, Q. Xu, N. Goyal, K. Singh, P. von Platen, Y. Saraf, J. Pino, A. Baevski, A. Conneau, and M. Auli, "XLS-R: Self-supervised Cross-lingual Speech Representation Learning at Scale," in *Proc. Interspeech 2022*, 2022, pp. 2278–2282.
- [18] Y.-T. Lee, X.-B. Chen, H.-S. Lee, J.-S. R. Jang, and H.-M. Wang, "Multi-task learning for acoustic modeling using articulatory attributes," in *2019 APSIPA ASC*, 2019, pp. 855–861.
- [19] K. Glocker and M. Georges, "Hierarchical multi-task transformers for crosslingual low resource phoneme recognition," in *Proc. 5th ICNLSP 2022*, Dec. 2022, pp. 187–192.
- [20] D. R. Mortensen, X. Li, P. Littell, A. Michaud, S. Rijhwani, A. Anastasopoulos, A. W. Black, F. Metze, and G. Neubig, "AlloVera: A multilingual allophone database," in *Proc. 12th LREC*, 2020, pp. 5329–5336.
- [21] X. Li, J. Li, F. Metze, and A. W. Black, "Hierarchical phone recognition with compositional phonetics," in *Interspeech 2021*. ISCA, Aug 2021, pp. 2461–2465.
- [22] S. Moran and D. McCloy, Eds., *PHOIBLE 2.0*. Jena: Max Planck Institute for the Science of Human History, 2019. [Online]. Available: <https://phoible.org/>
- [23] Q. Xu, A. Baevski, and M. Auli, "Simple and effective zero-shot cross-lingual phoneme recognition," *arXiv*, Sep 2021, arXiv:2109.11680.
- [24] A. Graves, S. Fernández, F. J. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," *Proc. 23rd ICML*, 2006.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [26] "eSpeak NG (version 1.51)," 2022. [Online]. Available: <https://github.com/espeak-ng/espeak-ng/tree/1.51>
- [27] H. Hammarström, R. Forkel, M. Haspelmath, and S. Bank, "glottolog/glottolog: Glottolog database 4.7," Dec. 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.7398962>
- [28] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," in *Proc. 12th LREC*, May 2020, pp. 4218–4222.
- [29] X. Li, D. R. Mortensen, F. Metze, and A. W. Black, "Multilingual phonetic dataset for low resource speech recognition," in *ICASSP 2021*, 2021, pp. 6958–6962. [Online]. Available: <https://github.com/xinjli/ucla-phonetic-corpus/releases/tag/v1.0>
- [30] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," *Advances in neural information processing systems*, vol. 32, 2019.
- [31] Y.-Y. Yang, M. Hira, Z. Ni, A. Astafurov, C. Chen, C. Puhersch, D. Pollack, D. Genzel, D. Greenberg, E. Z. Yang, J. Lian, J. Hwang, J. Chen, P. Goldsborough, S. Narenthiran, S. Watanabe, S. Chintala, and V. Quenneville-Bélaïr, "Torchaudio: Building blocks for audio and speech processing," in *ICASSP 2022*, 2022, pp. 6982–6986.
- [32] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush, "Transformers: State-of-the-art natural language processing," in *Proc. 2020 EMNLP*, 2020, pp. 38–45.
- [33] X. Li, F. Metze, D. R. Mortensen, A. W. Black, and S. Watanabe, "Phone inventories and recognition for every language," in *Proc. 13th LREC*, 2022, pp. 1061–1067.