

Characterization of Speech Similarity Between Australian Aboriginal and High-Resource Languages: A Case Study on Dharawal

Ting Dang*, Trini Manoj Jeyaseelan[†], Eliathamby Ambikairajah[†], Vidhyasaharan Sethu[†]

* The University of Melbourne, Australia

E-mail: ting.dang@unimelb.edu.au

[†] The University of New South Wales, Australia

E-mail: {trini_manoj.jeyaseelan, e.ambikairajah, v.sethu}@unsw.edu.au

Abstract—Australian Aboriginal languages are of significant cultural and linguistic value but remain severely underrepresented in modern speech AI systems. While state-of-the-art speech foundation models and automatic speech recognition excel in high-resource settings, they often struggle to generalize to low-resource languages, especially those lacking clean, annotated speech data. In this work, we collect and clean a speech dataset for Dharawal, a low-resource Australian Aboriginal language, by carefully sourcing and processing publicly available recordings. Using this dataset, we analyze the speech similarity between Dharawal and 107 high-resource languages using a pre-trained multilingual speech encoder. Our approach combines (1) misclassification rate analysis to assess language confusability, and (2) fine-grained similarity measurements using cosine similarity and Fréchet Inception Distance (FID) in the embedding space. Experimental results reveal that Dharawal shares strong speech similarity with languages such as Latin, Māori, Korean, Thai, and Welsh. These findings offer practical guidance for future transfer learning and model adaptation efforts, and underscore the importance of data collection and embedding-based analysis in supporting speech technologies for endangered language communities.

I. INTRODUCTION

Australia is home to a rich diversity of Aboriginal languages, many of which are critically endangered [1]. Among these, Dharawal is one of the Indigenous languages with deep cultural significance but limited linguistic resources [2]. Despite its historical and cultural importance, Dharawal, like many other Aboriginal languages, faces challenges in preservation and technological support due to its low-resource status.

The rapid advancement of AI technologies, particularly in speech processing, has led to significant breakthroughs in automatic speech recognition (ASR) [3], [4], language modeling [5], and multilingual understanding [6]. State-of-the-art models such as Whisper [7], Wav2Vec 2.0 [8], and VoxLingua107 ECAPA-TDNN [9] have demonstrated impressive capabilities in transcribing and understanding spoken language across diverse global languages. However, these advancements predominantly focus on high-resource languages, those with abundant annotated speech data, leaving low-resource languages, including Dharawal, largely underrepresented in modern AI-driven solutions **lowproject**. Moreover, the limited

availability of Dharawal speech recordings, especially the lack of transcriptions, further restricts the development of effective AI models for this language.

It is therefore essential to first collect Dharawal speech recordings for detailed analysis. While these recordings are largely unlabelled, a key initial step is to understand how Dharawal relates to high-resource languages within the multilingual speech representation space. By examining these relationships, we can more effectively leverage knowledge transfer and adaptation techniques. For instance, identifying languages with similar acoustic or phonetic characteristics enables the use of pretrained models from high-resource languages as a foundation for improving recognition or synthesis in low-resource languages. This strategy can substantially reduce dependency on large annotated low-resource languages datasets, which are often unavailable.

Existing studies have explored techniques such as transfer learning and adapters for other relatively low-resource languages [10]–[13] such as Georgian and Kazakh, but they rely on the availability of the transcribed speech recordings from publicly accessible datasets like Common Voice [14]. While there are increasing studies on Australian Aboriginal languages [15], there is currently no study on Dharawal, nor is there any publicly available labeled Dharawal speech dataset. Moreover, techniques commonly used for low-resource languages often rely on transcriptions [16] and therefore cannot be directly applied to untranscribed Dharawal speech recordings.

To bridge this gap, we first collected a Dharawal speech dataset. Further, we analyzed the speech similarity between Dharawal and other high resource languages, aiming to provide the first systematic quantification of speech similarity in the absence of transcriptions, to inform future transfer learning efforts. Specifically, we process the Dharawal speech recordings using pretrained multilingual encoders and quantify language similarities using two metrics: (1) the misclassification rate, which measures how frequently Dharawal is classified as other languages, and (2) cosine similarity and Fréchet Inception Distance (FID) scores [17], which quantify the similarity of the speech embeddings in the shared latent speech space. While the former provides a coarse-level estimate of confusion, the

latter offers a fine-grained, data-driven approach to measure affinities within the shared representational space. By quantifying these similarities, we can provide a foundation to better understand language relationships and guide the development of effective adaptation strategies for low-resource languages like Dharawal. The contributions of this paper are summarized as follows:

- We collected and cleaned the first speech dataset for the Australian Aboriginal language, Dharawal.
- We proposed the first speech embedding analysis for Dharawal by examining its misclassification rate and embedding similarity using a pretrained multilingual encoder, comparing it with other high-resource languages.
- Experimental results show that Dharawal is most similar to Latin, Māori, Korean, Thai, and Welsh. Although the ranking order differs slightly between misclassification rate and embedding similarity, the top 10 most similar languages largely overlap.

II. RELATED WORK

A. Dharawal

Dharawal is an Australian Aboriginal language spoken by people in the region surrounding modern day Sydney, New South Wales [18]. Like most Indigenous Australian languages, Dharawal faces challenges due to language shift and lack of intergenerational language transmission. Although efforts to revitalize and preserve the language are ongoing, available linguistic resources remain sparse. To the best of our knowledge, no prior work has collected the Dharawal speech recordings or analyzed speech embedding similarities using pre-trained speech encoders for Dharawal.

B. Language similarity identification

Language similarity detection plays a critical role in multilingual speech and language processing, particularly for low-resource languages. By identifying high-resource languages that are acoustically or structurally similar, researchers can leverage transfer learning [19], [20], cross-lingual adaptation [10], and multilingual modeling [6], [21] more effectively. Prior work has utilized phoneme-based comparisons [22], leveraging the linguistic similarities. However, the exploration of speech similarity beyond purely linguistic content remains limited. With the advancement of pretrained speech encoders, such as Wav2Vec 2.0 [8], which are trained on diverse languages and speech recordings, and the growing availability of large curated multilingual speech datasets like VoxLingua107 [23], it is now increasingly feasible to perform speech similarity identification without the need for extensive annotated data.

III. DATASET

The Dharawal speech dataset used in this study was derived from publicly available materials on the Dharawal Words website¹, which shares language resources from the native

¹<https://www.dharawalwords.com.au>

Dharawal speaker for educational and cultural preservation purposes. We collected the recordings following strict ethical guidelines, and the resulting dataset remains private and is used solely for research. The dataset included 475 utterances.

All recordings were manually reviewed to ensure clarity and audio quality. Specifically, each file was inspected to identify and remove segments containing excessive noise, corruption, or irrelevant content. Recordings with unclear speech or extended silence were excluded to maintain data integrity. This cleaning step helps minimize unnecessary padding and ensures that the resulting embeddings capture meaningful linguistic features rather than background noise.

The Dharawal speech segments were generally short and primarily consisted of individual words, with only a small number of utterances containing sentence-level content. To generate more meaningful samples for analysis, shorter recordings were concatenated in a way that preserved contextual coherence. This process was conducted using Audacity [24], an open-source audio editing tool, and included trimming silences at the beginning and end of each clip before assembling them into longer sequences. The resulting dataset comprises utterances approximately 10 to 15 seconds in length.

All audio files were saved in WAV format at a 16 kHz sampling rate. No data augmentation was applied in order to preserve the natural acoustic characteristics of the original recordings.

IV. METHODS

A. Overview

As shown in Figure 1, given a speech utterance x_i from our Dharawal dataset \mathcal{D} , we process the input through a pretrained multilingual encoder $f(\cdot)$ to obtain speech embeddings z_i . This enables us to perform a two-fold analysis: (i) estimating the misclassification rate by evaluating how often the embeddings of z_i are incorrectly classified as belonging to other high-resource languages, providing insight into the encoder’s discrimination capability for low-resource Aboriginal languages; and (ii) quantifying the embedding similarity between Dharawal and other languages using cosine similarity and FID scores, where cosine similarity captures the directional alignment between embedding vectors, and FID quantifies the overall statistical distance between their distributional properties (mean and covariance). This framework allows us to assess both the coarse-level separability and the fine-grained relative closeness of Dharawal embeddings within the multilingual embedding space.

B. Misclassification Rate Analysis for Language Identification

To identify high-resource languages similar to Dharawal, we employed a multilingual pretrained encoder that includes a wide variety of languages for analysis. Specifically, we used the state-of-the-art VoxLingua107-ECAPA-TDNN model [23], as they have been pretrained on a diverse corpus encompassing 107 distinct languages, ensuring broad linguistic representation across multiple language families. When presented with a

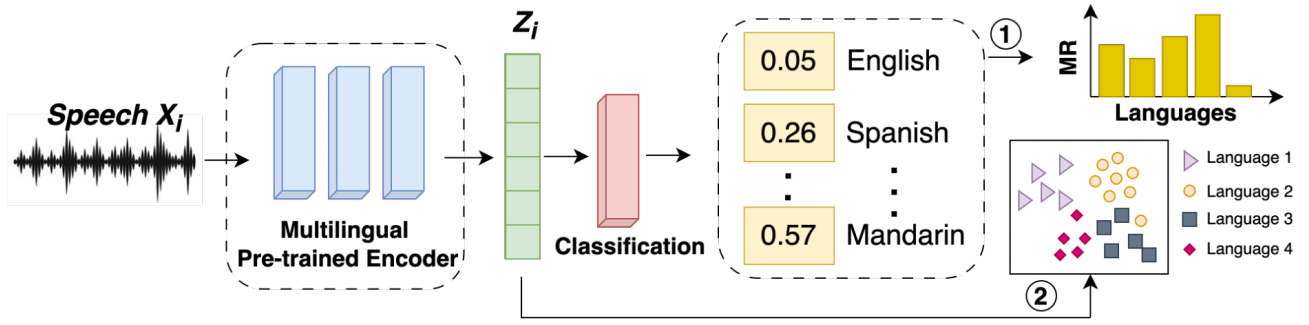


Fig. 1. System Overview. The input speech utterance in Dharawal, denoted as x_i , is processed by a multilingual pre-trained encoder to obtain embeddings z_i . These embeddings are then passed through a classification head to produce the final predicted language probabilities. We propose two approaches: (i) estimating the misclassification rate (MR) based on the probability outputs, and (ii) quantifying embedding similarity and FID scores in the latent space.

Dharawal speech utterance x_i , the model generates a probability distribution over the 107 language classes \hat{y}_i . The predicted label corresponds to the language with the highest assigned probability. The misclassification rate is determined by measuring the frequency with which Dharawal are incorrectly assigned to any of the known language categories as:

$$\text{MR} = \frac{1}{N} \sum_{i=1}^N I(\hat{y}_i \neq \text{Dharawal}) \quad (1)$$

where N is the total number of Dharawal utterances, \hat{y}_i is the predicted label for the i -th utterance, and $I(\cdot)$ is the indicator function that equals 1 if the condition is true, and 0 otherwise. A higher misclassification rate indicates greater similarity between Dharawal and the corresponding language, and vice versa.

C. Quantification of Embedding Similarity

To evaluate the embedding-level similarity between Dharawal and other high-resource languages, we employ cosine similarity and FID scores. These metrics quantify the difference between the speech embeddings of Dharawal and a target language in a shared embedding space.

a) Embedding Extraction: Speech embeddings are obtained using the encoder part of the pre-trained multilingual VoxLingua107-ECAPA-TDNN model [23]. Given a speech utterance x_i , the encoder maps it into a fixed-dimensional embedding vector:

$$z_i = f(x_i) \quad (2)$$

where z_i is the resulting speech embedding in a high-dimensional latent space as shown in Figure 1.

b) Cosine Similarity: For cosine similarity, the model pre-stores class centroids for the pretrained 107 high-resource languages, where μ_L denotes the centroid, i.e., the mean embedding vector of the L^{th} target language. We extract the centroid embedding to represent each high-resource language in the shared latent space. For the Dharawal input utterances x_i , we compute their corresponding embeddings and calculate

the mean to obtain the centroid representation μ_{Dharawal} as:

$$\mu_{\text{Dharawal}} = \frac{1}{N} \sum_{i=1}^N z_i \quad (3)$$

The cosine similarity between Dharawal and a target language L is then computed as:

$$\text{CosSim}(\mu_{\text{Dharawal}}, \mu_L) = \frac{\mu_{\text{Dharawal}} \cdot \mu_L}{\|\mu_{\text{Dharawal}}\| \|\mu_L\|} \quad (4)$$

A higher cosine similarity score indicates that the embeddings of Dharawal are statistically closer to those of the target language, suggesting stronger linguistic or acoustic similarities. Conversely, a lower cosine similarity implies greater divergence in the representation space.

c) FID scores: Regarding FID scores, the computation requires estimating the distributional properties of the embedding space for each language, specifically the mean and covariance of the embeddings. Therefore, we need to collect speech recordings for each high-resource languages to obtain representative embedding clusters. For comparative analysis, we select the top 10 most similar high-resource languages, based on prior misclassifications rates, as the target languages for FID evaluation. The final dataset for FID score analysis includes the 10 selected high-resource languages, alongside the Dharawal speech recordings.

Specifically, let $(\mu_{\text{Dharawal}}, \Sigma_{\text{Dharawal}})$ and (μ_L, Σ_L) be the mean vectors and covariance matrices of the embeddings for Dharawal and the target language, respectively. The FID score are then calculated as follows:

$$\text{FID}(\mathcal{Z}_D, \mathcal{Z}_L) = \|\mu_D - \mu_L\|^2 + \text{Tr}(\Sigma_D + \Sigma_L - 2(\Sigma_D \Sigma_L)^{\frac{1}{2}}) \quad (5)$$

where $\|\cdot\|^2$ denotes the squared Euclidean distance, and $\text{Tr}(\cdot)$ represents the matrix trace operation. A lower FID score indicates that the embeddings of Dharawal are statistically closer to those of the target language, and vice versa.

V. EXPERIMENTAL SETUP

A. Multilingual pre-trained encoder

The multilingual pre-trained encoder VoxLingua107-EPACA-TDNN is built on the ECAPA-TDNN (Emphasized



Fig. 2. Multilingual pre-trained encoder architecture for similarity characterization.




Fig. 3. Misclassification rates for the top 10 languages identified as similar to Dharawal using VoxLingua107-EPACA-TDNN, respectively.

Channel Attention, Propagation, and Aggregation Time Delay Neural Network) architecture [25], as shown in Figure 2. It was trained on the VoxLingua107 dataset [23], which consists of over 6,000 hours of speech from 107 different languages. Specifically, it first extracts frame-level features from raw audio inputs using a series of 1D convolutional layers, which capture local acoustic patterns and temporal dependencies directly from the waveform. To enhance feature representation, Squeeze-Excitation (SE) blocks are employed, dynamically recalibrating channel-wise responses to emphasize informative channels. Temporal processing is managed by TDNN layers, which apply 1D convolutions over varying contexts to incorporate long-range dependencies. An attention mechanism then aggregates frame-level features, weighting each frame’s contribution based on its importance, and outputs a 256-dimensional embedding vector that encapsulates the global linguistic and acoustic characteristics of the speech segment, optimized for language classification and similarity analysis.

B. Implementation details

During the inference phase, the multilingual encoder processes input Dharawal utterances to evaluate (i) the misclassification rate and (ii) embedding similarity. For the misclassification rate, we computed the error using the softmax output probabilities. For cosine similarity and FID scores, we extracted the 256-dimensional embeddings for each of the 107 pretrained languages and compared it to the embeddings of Dharawal to quantify similarity. To ensure a fair comparison, the number of utterances used for FID computation was kept consistent across all languages.

All experiments were conducted using Python 3.11.12 with PyTorch and the SpeechBrain libraries. All computations were performed on Google Colab using CPU resources.

VI. RESULTS

A. Misclassification rate

The misclassification rate for the pretrained encoder are shown in Figure 3. It displays the top 10 languages with the highest misclassification rates for Dharawal. It is evident that Dharawal is most frequently misclassified as Latin, with a misclassification rate of 14.52%, indicating a notable degree of similarity between the two. Classical Latin has a relatively simple and consistent vowel system and syllable structure (often Consonant-Vowel or Consonant-Vowel-Consonant) [26], characteristics that are also common in many Australian Aboriginal languages. Furthermore, both may lack certain consonant clusters or tonal variation, making them acoustically similar at a superficial level. This is followed by Māori, Korean, Thai, and Welsh, which also exhibit relatively high confusion rates. The confusion with Māori is more expected. Both languages are indigenous to the southern hemisphere and are known for having small phoneme inventories, frequent use of open syllables (Consonant-Vowel patterns), and clear, rhythmic prosody [27]. These shared features can easily cause acoustic overlap in models that rely on phonetic features alone. These similarities suggest opportunities to leverage the acoustic patterns of similar languages for data augmentation or transfer learning, to further improve AI models for under-resourced Dharawal languages.

B. Embedding similarity

To provide a more fine-grained analysis of embedding similarities, we report cosine similarity and FID scores between the encoders and the top 10 most similar high-resource languages.

For cosine similarity and FID scores of the top 10 most similar languages, as shown in Table I and Figure 4, the same set of languages appears with nearly identical ranking orders. However, these fine-grained similarity measures show slight differences in ranking when compared to the misclassification rates. Notably, the embeddings show the highest cosine similarity and FID scores for Thai, Welsh, and Korean, while Latin and Māori which are ranked highest in the misclassification rate are ranked least using the cosine similarity and FID scores. This discrepancy arises because cosine similarity, FID score and misclassification rate capture fundamentally different aspects of language similarity and model behavior. Cosine similarity and FID score measure the geometric closeness of language embeddings in the learned feature space, reflecting intrinsic representational similarity based on shared phonetic, phonological, or acoustic patterns. Thus, languages like Thai, Welsh, and Korean may cluster closely due to underlying phonetic traits that the model encodes similarly. In contrast, the misclassification rate reflects practical classification challenges

TABLE I
COSINE SIMILARITY BETWEEN DHARABAL SPEECH EMBEDDINGS AND SPEECH EMBEDDINGS FROM OTHER 10 MOST SIMILAR HIGH-RESOURCE LANGUAGES.

Languages	Thai	Welsh	Korean	Tibetan	Estonian	English	Javanese	Catalan	Māori	Latin
Cos similarity	0.9078	0.8969	0.8937	0.8845	0.8785	0.8749	0.8743	0.8562	0.8533	0.8238




Fig. 4. FID scores between Dharawal speech embeddings and the speech embeddings of the 10 most similar languages.

and model errors when distinguishing between languages. The high misclassification rates for Latin and Māori suggest that the model struggles to distinguish these languages from others. This may be due to partially overlapping acoustic features, insufficient training data, or noise and variability within the Māori recordings. Despite their relatively low overall cosine similarity, the distribution or clustering of these languages may contain subsets that overlap significantly with others, leading to confusion during classification. These metrics complement each other, offering a comprehensive picture of language similarity and model performance.

C. Visualization

To further explore the embedding space, we visualize the latent representations of the top 10 most similar languages with respect to Dharawal. The t-SNE (t-distributed Stochastic Neighbor Embedding) plot in Figure 5 demonstrates patterns that largely align with the computed similarity metrics. Dharawal exhibits the closest proximity to Thai (cyan) and Korean (light orange) in the latent space, consistent with their high computed similarity scores.

For Javanese (pink) and Estonian (light green), the main cluster position closely to Dharawal, but scattered outlier points distributed across the latent space, likely reflecting significant internal language variability, cause the centroid to shift away from Dharawal, resulting in moderate similarity ranking among the ten languages. Catalan (purple) is positioned far from Dharawal, therefore ranked the 8th in the top 10 languages. Notably, small clusters of Māori (light purple) appear on the right side of the Dharawal clusters, in close proximity to them, which may explain the observed high misclassification rates between these languages. However,




Fig. 5. t-SNE plot of the embeddings for Dharawal and other high-resource languages.

sparse Māori data points and large clusters extending toward the left side of the embedding space pull the Māori centroid away from Dharawal, consequently reducing their computed similarity score. These observations suggest that Māori and Dharawal share substantial underlying similarity in the latent representation in general.

VII. CONCLUSION

In this work, we collect and clean a dataset for the Dharawal language, contributing a valuable resource to the growing body of speech data for Australian Aboriginal languages. We investigated two approaches for analyzing speech similarity using pre-trained multilingual speech encoders. Specifically, we evaluated classification separability via misclassification rates and quantified embedding similarity through cosine similarity and Fréchet Inception Distance scores.

Our results consistently indicate that Dharawal exhibits strong acoustic similarity to high-resource languages such as Latin, Māori, Korean, Thai, and Welsh. This finding suggests that speech models trained on these languages may serve as effective bases for transfer learning or adaptation to support Dharawal and potentially other low-resource Aboriginal languages—without requiring large-scale transcribed data.

This study underscores the importance of systematic similarity analysis for informing downstream development of speech technologies for endangered languages. Future work could expand on these insights by exploring multiple model architectures and larger datasets as they become available, developing lightweight adaptation techniques based on the most acoustically similar high-resource languages, and applying this framework to additional Indigenous Australian languages to further support linguistic preservation and accessibility.

ACKNOWLEDGMENT

The authors would like to thank the School of Electrical Engineering and Telecommunications at UNSW Sydney, Australia, for providing funding for this research initiative.

REFERENCES

- [1] J. Simpson and G. Wigglesworth, "Language diversity in indigenous australia in the 21st century," *Current Issues in Language Planning*, vol. 20, no. 1, pp. 67–80, 2019.
- [2] A. Capell, "Aboriginal languages in the south central coast, new south wales: Fresh discoveries," *Oceania*, vol. 41, no. 1, pp. 20–27, 1970.
- [3] H. Kheddar, M. Hemis, and Y. Himeur, "Automatic speech recognition using advanced deep learning approaches: A survey," *Information Fusion*, p. 102422, 2024.
- [4] Y. Chu, J. Xu, Q. Yang, *et al.*, "Qwen2-audio technical report," *arXiv preprint arXiv:2407.10759*, 2024.
- [5] Y. Fathullah, C. Wu, E. Lakomkin, *et al.*, "Prompting large language models with speech recognition abilities," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2024, pp. 13 351–13 355.
- [6] V. Pratap, A. Tjandra, B. Shi, *et al.*, "Scaling speech technology to 1,000+ languages," *Journal of Machine Learning Research*, vol. 25, no. 97, pp. 1–52, 2024.
- [7] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *International conference on machine learning*, PMLR, 2023, pp. 28 492–28 518.
- [8] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "Wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [9] M. Chaitra, A. Mandal, and S. Mukherjee, "Study of ecapa-tdnn models for spoken language identification task," in *2023 IEEE World Conference on Applied Intelligence and Computing (AIC)*, IEEE, 2023, pp. 233–237.
- [10] T. Tu, Y.-J. Chen, C.-c. Yeh, and H.-Y. Lee, "End-to-end text-to-speech for low-resource languages by cross-lingual transfer learning," *arXiv preprint arXiv:1904.06508*, 2019.
- [11] A. Romanenko, "Robust speech recognition for low-resource languages," Ph.D. dissertation, Universität Ulm, 2022.
- [12] W. Hou, H. Zhu, Y. Wang, *et al.*, "Exploiting adapters for cross-lingual low-resource speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 317–329, 2021.
- [13] M. Kim, K. Jang, and H. Kim, "Improving cross-lingual phonetic representation of low-resource languages through language similarity analysis," in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2025, pp. 1–5.
- [14] R. Ardila, M. Branson, K. Davis, *et al.*, "Common voice: A massively-multilingual speech corpus," *arXiv preprint arXiv:1912.06670*, 2019.
- [15] E. Ambikairajah, J. Wu, T. Dang, and V. Sethu, "A study of speech embedding similarities between australian aboriginal and high-resource languages," in *Proc. Interspeech 2025*, 2025, pp. 1498–1502.
- [16] T. Ahmed, M. Suffian, M. Y. Khan, and A. Bogliolo, "Discovering lexical similarity using articulatory feature-based phonetic edit distance," *IEEE Access*, vol. 10, pp. 1533–1544, 2021.
- [17] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "Gans trained by a two time-scale update rule converge to a local nash equilibrium," *Advances in neural information processing systems*, vol. 30, 2017.
- [18] D. E. A. B. Timbery-Beller, "Dharawal,"
- [19] J. Yi, J. Tao, Z. Wen, and Y. Bai, "Language-adversarial transfer learning for low-resource speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 3, pp. 621–630, 2018.
- [20] Z. Byambadorj, R. Nishimura, A. Ayush, K. Ohta, and N. Kitaoka, "Text-to-speech system for low-resource language using cross-lingual transfer learning and data augmentation," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2021, no. 1, p. 42, 2021.
- [21] S. Toshiwal, T. N. Sainath, R. J. Weiss, *et al.*, "Multilingual speech recognition with a single end-to-end model," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, IEEE, 2018, pp. 4904–4908.
- [22] D. Ingram and E. Babatsouli, "Cross-linguistic phonological acquisition," *The Handbook of Clinical Linguistics, Second Edition*, pp. 407–419, 2024.
- [23] J. Valk and T. Alumäe, "VoxLingua107: A dataset for spoken language recognition," in *Proc. IEEE SLT Workshop*, 2021.
- [24] T. Audacity, "Audacity," *The name audacity (R) is a registered trademark of dominic mazzoni retrieved from http://audacity.sourceforge.net*, 2017.
- [25] S. He, H. Zhang, W. Rao, *et al.*, "Hierarchical speaker representation for target speaker extraction," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2024, pp. 10 361–10 365.
- [26] A. Cser, "Aspects of the phonology and morphology of classical latin," Ph.D. dissertation, Pázmány Péter Katolikus Egyetem, 2016.
- [27] R. Harlow, *Maori: A linguistic introduction*. Cambridge University Press, 2007.