



Speaker- and Text-Independent Estimation of Articulatory Movements and Phoneme Alignments from Speech

Tobias Weise^{1,2}, Philipp Klumpp², Kubilay Can Demir¹, Paula Andrea Pérez-Toro^{2,4}, Maria Schuster³, Elmar Noeth², Bjoern Heismann², Andreas Maier², Seung Hee Yang¹

¹Speech & Language Processing Lab. Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany

²Pattern Recognition Lab. Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany

³Department of Otorhinolaryngology, Head and Neck Surgery. Ludwig-Maximilians University, Munich, Germany

⁴GITA Lab, Faculty of Engineering. Universidad de Antioquia, Medellín, Colombia

tobias.weise@fau.de

Abstract

This paper introduces a novel combination of two tasks, previously treated separately: acoustic-to-articulatory speech inversion (AAI) and phoneme-to-articulatory (PTA) motion estimation. We refer to this joint task as acoustic phoneme-to-articulatory speech inversion (APTAI) and explore two different approaches, both working speaker- and text-independently during inference. We use a multi-task learning setup, with the end-to-end goal of taking raw speech as input and estimating the corresponding articulatory movements, phoneme sequence, and phoneme alignment. While both proposed approaches share these same requirements, they differ in their way of achieving phoneme-related predictions: one is based on frame classification, the other on a two-staged training procedure and forced alignment. We reach competitive performance of 0.73 mean correlation for the AAI task and achieve up to approximately 87% frame overlap compared to a state-of-the-art text-dependent phoneme force aligner.

Index Terms: speech inversion, attention, phoneme alignment, wav2vec 2.0, HPRC, tract variables, multi-task learning

1. Introduction

In phonetics, articulatory configurations are analyzed to understand how different sounds are produced and how they can be classified into phonemes within a particular language's phonological system. Articulators refer to the various parts of the vocal tract and other structures (e.g. tongue, lips, palate) involved in the production of sounds. They are typically measured by placing sensor coils, in a procedure called electromagnetic articulography (EMA), and tracking the position and movement over time during speech. These sensor coordinates are naturally speaker-specific since they depend on the particular vocal tract anatomy of the recorded speaker. Tract Variables (TVs), introduced by Brownman et. al. [1], on the other hand, combine multiple individual vocal tract articulator movements, that achieve a specific linguistic objective, into defined *gestures* relevant to articulation. Transformations were introduced by Ji [2] to convert EMA sensor coordinates into TVs, which were shown to be less speaker dependent [3] than the original measurements.

The problem of inverting an original speech signal back to its articulator positions is referred to as acoustic-to-articulatory speech inversion (AAI), which can involve TVs or EMA coordinates as targets. This task has been studied speaker-dependent and speaker-independently in literature: multi-task learning (MTL) [6, 7], generative adversarial networks [8], the appli-




Figure 1: Nine tract variables (TVs), used for speaker-independent articulatory speech inversion. Adapted from [4, 5].




Figure 2: Proposed APTAI model, based on wav2vec2.0 fine-tuning via frame-classification and TV regression.

cation to dysarthric speech [9], and speech therapy [10, 11, 12], the incorporation of fundamental frequency [13], and others [14, 15, 16] have been explored. A related but less studied problem is taking a sequence of phonemes and mapping it to articulator movements (PTA): gated bidirectional recurrent neural networks [17], attempts to model the entire vocal tract [18], comparative studies [19], and feed-forward transformers [20] have been applied, where the latter authors also applied it to AAI in a speaker-dependent setting.

Phoneme recognition can be described as taking an audio signal as input and producing the corresponding frame-asynchronous phoneme sequence. However, the frame-synchronous relation [21] is required for the task of phoneme alignment [22, 23, 24], boundary detection, and segmentation [25]. This paper focuses on phoneme recognition and subsequent alignment to the individual frames, which can be beneficial e.g. during speech therapy [26, 27]. Here, we explore frame-wise classification and forced alignment. Our upper bound is a state-of-the-art (SOTA) text-dependent force aligner. This system relies on both audio and transcriptions as input, which are converted from graphemes to phonemes.

This paper introduces APTAI, a novel combination of AAI and PTA in combination with phoneme recognition and alignment. We require that resulting models predict end-to-end (in a therapeutic context) while working speaker- and text-independently during inference. To this end, two different approaches are explored, with Figure 1 illustrating the TV regression targets to model articulation.

2. Proposed Approach(es)

This paper introduces two approaches, sharing the same requirements outlined in the last paragraph of the introduction. Both make use of MTL optimization, composed of articulator movement regression and phoneme prediction paired with alignment. The main difference is the way they deal with the phoneme-related objective: APTAI is based on frame classification, whereas f-APTAI utilizes forced alignment during a two-staged training procedure. Our code is available online¹.

Both approaches make use of self-supervised learning (SSL) models but in different setups. Taking ASR as an example, SOTA performance has been achieved using this paradigm, which includes pre-training on large amounts of unlabeled data and fine-tuning on a smaller, labeled dataset relevant to the desired downstream task. We chose wav2vec2 [28], which optimizes a contrastive loss during pre-training to learn a finite set of speech representations. These can be fine-tuned for a broad set of applications, with ASR as the original intended use case. Thus, such embeddings are expected to capture meaningful features of speech that are relevant for phonemes, which in turn can be identified by specific articulator configurations.

Table 1: *Fine-tuned phoneme recognizer results (PER [%] ↓), using CP train/dev splits, for different pre-trained models.*

wav2vec2-	CP-test	HPRC-N	HPRC-F
base-960h	17.77	10.10	19.98
large-960h	18.71	11.47	24.27
large-1v60	9.75	4.96	13.76
large-960h-1v60	9.30	4.55	10.69
large-robust	8.83	4.45	10.53
xls-r-300m	11.70	7.77	19.38
xls-r-1b	18.50	12.69	27.29
large-xlsr-53	10.17	5.40	14.55

2.1. Frame Classification: APTAI

Of the two proposed approaches, APTAI follows a more classical setup, refer to Figure 2 for an overview. The general idea is to fine-tune wav2vec2 to make use of its pre-trained speech representations, which is the reason why we keep the feature extractor frozen (pre-trained weights), and only train the transformer layers (pre-trained initialization) in addition to two added heads (randomly initialized). Furthermore, we add a convolutional layer (fixed parameters), which behaves like a low-pass (sinc) filter, adapted from [29]. This enforces the smoothness of the predicted TV trajectories, which is required since frame-based signal regression typically suffers from high-frequency noise between the individual frame predictions.

An 16kHz input speech signal $x(t)$ is divided into T frames $\mathbf{x}_t \in \mathbb{R}^{512}$ at 49 Hz by the feature encoder. After passing the transformer layers, producing $\mathbf{h}_t \in \mathbb{R}^{1024}$, the TV head takes this output and ultimately predicts $\hat{\mathbf{y}}_t^{tv} \in \mathbb{R}^{TV}$ smoothed $TV = 9$ values for each frame t . As part of the MTL goal, this head optimizes the reconstruction mean square error (MSE)

loss between the predicted $\hat{\mathbf{y}}_t^{tv}$ and ground truth \mathbf{y}_t^{tv} TV values, which is expressed in the second term of Equation 1. The phoneme head also takes \mathbf{h}_t as input and predicts a probability distribution $\hat{p}_{t,c}$ over $C = 45$ phoneme labels per frame t , with $c \in C$. This frame-wise classification is optimized via cross-entropy (CE) loss between the predicted $\hat{p}_{t,c}$ and ground truth $p_{t,c}$ probability distribution (see first term in Equation 1). Applying *softmax* to the resulting logits and choosing the phoneme label c that yields the maximum probability per frame t will result in an alignment, whilst a phoneme sequence can be obtained by grouping over the individual frame predictions. Finally, Equation 1 shows the MTL loss \mathcal{L}_{FC} for the APTAI approach, with λ as weighting factor.

$$\mathcal{L}_{FC} = -\frac{1}{T} \sum_{t=1}^T \sum_{c=1}^C p_{t,c} \log(\hat{p}_{t,c}) + \lambda \frac{1}{T} \sum_{t=1}^T (\mathbf{y}_t^{tv} - \hat{\mathbf{y}}_t^{tv})^2 \quad (1)$$

2.2. Forced Alignment: f-APTAI

The idea behind the second approach f-APTAI is to make use of hidden representations from a fine-tuned phoneme recognizer in combination with a forced alignment of the predicted output phoneme sequence. To this end, we use a two-staged approach during training, depicted in Figure 3. We make use of different datasets for the two stages, more details in section 3.1.

For the first stage, we fine-tune the same SSL architecture (wav2vec2) used in APTAI, by adding a linear layer producing $\mathbf{l}_t \in \mathbb{R}^{C_0}$ representing the same $C = 45$ phoneme labels with the addition of a blank token \emptyset , per frame $t \in T$ (see section 2.1). Similar to the ASR application, we optimize this model using the connectionist temporal classification (CTC) loss. This optimization behaves like a state machine, similar to hidden markov models (HMM), and only requires a phoneme sequence as additional input during training. However, CTC does not produce an alignment but rather outputs a frame-asynchronous (in our case) phoneme label sequence through a frame-synchronous decoding procedure (beam search), utilizing the blank token and multiple possible alignment paths. Given a true phoneme label sequence \mathcal{W} , then \mathcal{S} represents all possible paths that map from \mathcal{W} to T by removing repeated labels and blanks. Then, $P(s_t | \mathbf{l}_t)$ represents the output of the model at t by applying *softmax* to \mathbf{l}_t , with $[s_{1:T}] \in \mathcal{S}$. Adapted from [21], the CTC loss can be defined as:

$$\mathcal{L}_{CTC} = -\log \sum_{\mathcal{S}} \prod_{t=1}^T P(s_t | \mathbf{l}_t) \quad (2)$$

The second stage of f-APTAI incorporates the frozen model trained during stage-1. Specifically, two parts are extracted and used during training of stage-2: the predicted CTC-based phoneme sequence (upper bound for stage-2) and the output of the last transformer layer. Here, let the former be $[p_{1:N}] \in \mathcal{P}^N$, where $p_n \in C$, and N the maximum sequence length. The last transformer layer output can be expressed as matrix \mathbf{H} , consisting of $\mathbf{h}_t \in \mathbb{R}^{1024}$ column vectors, with $t \in T$. This can be understood as acoustic phoneme embeddings since the stage-1 objective (see Equation 2) led to accordingly optimized weights. A principal component analysis (PCA) of these embeddings (extracted from the HPRC-N dataset, see section 3.1) can be seen in Figure 4. The setup is similar to [30] and shows good speaker independence with phoneme clustering of exemplary chosen elongated vowels, a fricative, nasal, and plosive. The performed neural forced alignment is inspired by [23] and has the goal of producing a monotonic alignment, such that it aligns each phoneme label p_n to a

¹<https://github.com/tobwei/APTAI>

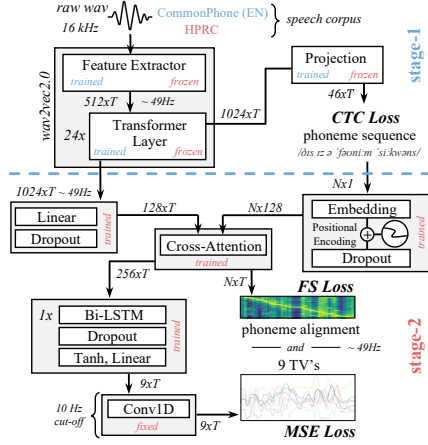


Figure 3: Proposed f -APTAI model, based on TV regression and a two-staged forced alignment via cross-attention.

subset of consecutive hidden frame representations \mathbf{h}_t . Therefore, one of the MTL optimization goals of f -APTAI is to learn a matrix $\mathbf{A} \in \mathbb{R}^{N \times T}$ that aligns \mathcal{P}^N to \mathbf{H} . This objective is centered around a cross-attention computation between a learned linear projection of \mathbf{h}_t to $\mathbf{h}_t^p \in \mathbb{R}^{128}$ resulting in $\mathbf{H}_p \in \mathbb{R}^{T \times 128}$, and a learned embedding of \mathcal{P}^N . This embedding is created via projection of each $p_n \in \mathcal{P}^N$ to \mathbb{R}^{128} and the addition of a sinusoidal positional encoding [31], ultimately resulting in matrix $\mathbf{P} \in \mathbb{R}^{128 \times N}$. Finally, the cross-attention layer computes the alignment matrix $\mathbf{A} = \text{softmax}(\mathbf{H}_p \cdot \mathbf{P})$. We constrain \mathbf{A} to be monotonic and diagonal, which is inspired by the forward-sum (FS) loss used in HMM systems, and adapted from [22, 24]. See the first term in Equation 3, where \mathcal{O} is the optimal alignment.

$$\mathcal{L}_{FA} = - \sum_{\mathbf{H}_p, \mathbf{P} \in \mathcal{O}} \log P(\mathbf{P} | \mathbf{H}_p) + \lambda \frac{1}{T} \sum_{t=1}^T (\mathbf{y}_t^{tw} - \hat{\mathbf{y}}_t^{tw})^2 \quad (3)$$

Additionally, the cross-attention layer produces a hidden representation matrix $\in \mathbb{R}^{256 \times T}$. This sequence of column vectors over T frames serves as input for the TV regression part of the f -APTAI model. Initially, it is passed through a single bi-directional long short-term memory (LSTM) layer, the output of which is ultimately projected to \mathbb{R}^{TV} . Moreover, the same fixed-parameter convolutional low-pass (sinc) filter as in APTAI is used to ensure the prediction of smooth TV trajectories $\hat{\mathbf{y}}_t^{tw}$. Consequently, the same MSE loss is also optimized, see the second term in Equation 3.

3. Experimental Setup

It should be noted that our upper bound for both approaches, in terms of phoneme recognition and alignment, is a SOTA [23] text-dependent force aligner from WebMAUS [32]. The reason for this is that we produce our ground truth phoneme labels and time steps via this web API. We make use of CommonPhone (see section 3.1) for its robustness and this dataset utilized the same process, so we apply the same to HPRC, the second dataset that we use to guarantee compatibility.

3.1. Datasets

One of the two datasets that we use during experiments is Common Phone (CP) [33], which is based on the crowd-sourced Common Voice [34]. Here, we utilize the English subset (45

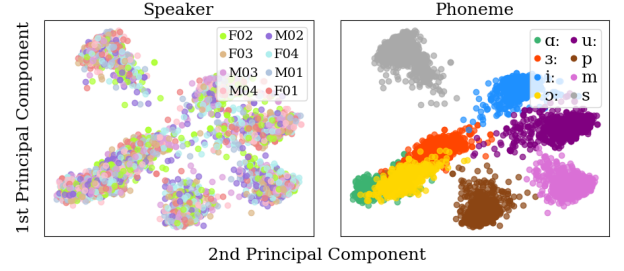


Figure 4: PCA of the embeddings from the best-performing fine-tuned phoneme recognition model from Table 1.

Table 2: Leave-one-speaker-out results (mean and deviation across eight test speakers) for the two proposed approaches.

Model, Test Data	PCC \uparrow	RSME[<i>mm</i>] \downarrow	PER[%] \downarrow	Overlap[%] \uparrow
APTAI, HPRC-N	0.73 \pm 0.03	0.67 \pm 0.03	6.25 \pm 1.30	87.38 \pm 1.16
APTAI, HPRC-F	0.69 \pm 0.03	0.72 \pm 0.03	6.41 \pm 1.76	84.91 \pm 1.93
f -APTAI, HPRC-N	0.71 \pm 0.03	0.68 \pm 0.03	4.36 \pm 0.07	76.18 \pm 1.59
f -APTAI, HPRC-F	0.65 \pm 0.03	0.74 \pm 0.03	10.29 \pm 3.62	72.93 \pm 2.92

phoneme labels). The main motivation behind using CP is that we want to build a robust system. When comparing CP to e.g. TIMIT [35], this robustness becomes evident: one is recorded in the same acoustically controlled environment with professional equipment, and the other is based on recordings from people’s smartphones in many different uncontrolled environments.

The second dataset we use contains articulator-related information in the form of EMA sensor data. This dataset is the Haskins Production Rate Comparison (HPRC) [36], which contains recordings from four female and four male subjects reciting 720 phonetically balanced IEEE sentences at “normal” (HPRC-N) and “fast” (HPRC-F) speaking rates. The speakers in this dataset repeat utterances, however, we randomly select only one repetition per utterance and speaker. Furthermore, we used the MAUS aligner to create our ground truth phoneme labels and time steps. This dataset comes with labels from another aligner, but we wanted to make it compatible with the CP dataset. Next, we performed pre-processing on the EMA data: some of the coordinates contained *NaN* values, where we applied linear interpolation to remedy this problem before low-pass (Butterworth) filtering the sensor data with 20 Hz to eliminate recording related noise. After this, the EMA coordinates were transformed into nine TVs (see Figure 1) and some final processing was applied to them. The original EMA data was sampled at 100 Hz, resulting in TVs at the same rate. We resampled them to 49 Hz to synchronize them with the output frame rate of wav2vec2. Finally, we applied utterance-wise z-score normalization based on the individual TVs.

3.2. Model Evaluation

We evaluate the APTAI task in terms of the two MTL sub-objectives. The articulation regression performance is evaluated using two well-known metrics: the root mean square error (RMSE) based on the normalized values and the Pearson correlation coefficient (PCC). To evaluate the phoneme recognition and alignment performance, we use the phoneme error rate (PER), where the ground truth is based on the webMAUS grapheme-to-phoneme conversion. Phoneme alignment is also evaluated regarding this text-dependent upper bound, using the frame-wise overlap (percentage of correctly predicted frames).




Figure 5: Example model prediction (APTAI) and ground truth for an unseen speaker (only a selection of all TVs is shown).

3.3. Model Training

The following setup was used to train/validate our two proposed approaches, using the PyTorch framework. For CP, we used the official train/dev/test splits. To test the performance of our models, we used HPRC. Here, we applied leave-one-speaker-out testing, i.e., data from seven speakers was used for training/validation (90%/10%), and the data of the remaining speaker was used to test (separated by speaking rates). Additionally, we performed the training split in such a way that only unseen utterances were used for validation. The same optimizer (Adam), learning rate ($1e-5$), learning-rate scheduler (warm-up, static, and decaying epochs), batch size of 5, and model selection metric (TV RMSE) were used for both proposed approaches. We experimented with MTL strategies (e.g. alternating epochs) but with no improvement in performance.

APTAI, utilizing `wav2vec2-large-robust` (see Table 1), was trained for 20 epochs, with 20% dropout, and combined HPRC-N and -F for training/validation. In terms of the MTL loss optimization, we set $\lambda = 1$ thus weighting both tasks equally, which resulted in the best performance.

Fine-tuning of the phoneme recognizer for stage-1 of f -APTAI was based on `wav2vec2-large-robust` (best performance, see Table 1) with a batch size of 2, 160 epochs, learning rate of $5e-6$, a final dropout of 10%, and model selection based on validation PER. For stage-2, we trained for 60 epochs, used only HPRC-N (since including F would negatively impact the PER of stage-1), set $\lambda = 0.4$, and $N = 60$, with shorter phoneme sequences being padded. Finally, the implementation of the FS loss was taken from [24].

4. Results and Discussion

Table 1 reveals that CP is a noisy dataset, while HPRC is not. This results in better PER for "normal" speaking rates, while "fast" are more challenging (also for human listeners), with `wav2vec2-large-robust` performing best.

Table 2 shows the main evaluation test results of the introduced APTAI task, conducted in a speaker-independent (LOSO) setting. Figure 5 illustrates prediction performance, showing a selection of TVs for improved readability, whilst Figure 6 shows all TVs individually. In terms of TV metrics, both models perform similarly, with APTAI achieving the best mean PCC of 0.73. Comparing this result to other works is difficult since setups are not uniform (e.g. trimming of silence), and reproduced results do not match originally reported ones [6, 16]. However, reported speaker-independent PCC results on HPRC roughly range from 35% to 76%, so we achieve competitive




Figure 6: Ground truth and model prediction (APTAI) in red color, of an unseen speaker (refer to Figure 1 as legend).

Table 3: Individual TV metrics, in terms of mean and deviation across the leave-one-speaker-out experiments (APTAI model).

TV's	HPRC-N		HPRC-F	
	PCC \uparrow	RSME[mm] \downarrow	PCC \uparrow	RSME[mm] \downarrow
LA	0.87 \pm 0.03	0.49 \pm 0.06	81.76 \pm 4.89	0.57 \pm 0.07
LP	0.75 \pm 0.08	0.66 \pm 0.10	66.93 \pm 8.57	0.75 \pm 0.10
JA	0.82 \pm 0.04	0.57 \pm 0.06	73.97 \pm 4.19	0.67 \pm 0.06
TTCL	0.84 \pm 0.04	0.54 \pm 0.06	81.85 \pm 3.25	0.56 \pm 0.05
TTCD	0.79 \pm 0.04	0.61 \pm 0.06	74.14 \pm 5.48	0.67 \pm 0.06
TMCL	0.82 \pm 0.03	0.57 \pm 0.04	79.38 \pm 2.47	0.60 \pm 0.04
TMCD	0.37 \pm 0.11	1.07 \pm 0.09	27.94 \pm 11.34	1.13 \pm 0.09
TBCL	0.77 \pm 0.04	0.64 \pm 0.05	74.36 \pm 4.36	0.67 \pm 0.06
TBCD	0.54 \pm 0.15	0.88 \pm 0.14	56.57 \pm 14.53	0.85 \pm 0.14

performance. In terms of phoneme recognition and alignment, frame classification outperforms the forced alignment approach by 11.20%, achieving a frame overlap of 87.38%. Shih et. al. [24] reported that in their experiments, a wider receptive field lead to alignment instability. The fact that we use hidden transformer representations, capturing weighted global sequence dependencies, might explain the reduced alignment performance, which requires future research. Overall, the work of Siriwardena et. al. [7] is similar, however, they report a PER of approx. 27% (and no alignment metric) since they see the phoneme-related objective as an auxiliary task to improve TV-related performance, while we see both tasks as equally important.

When looking at Table 3 and Figure 6, it is noticeable that especially the regression of TMCD and TBCD perform significantly worse when compared to the other TVs, hampering the overall mean PCC. This needs further investigation since other papers do not seem to suffer from this problem.

5. Conclusion

This paper introduced APTAI, a novel combination of two tasks previously viewed separately. We investigated two different approaches, sharing the same robust requirements but differing mainly in their method of phoneme prediction and alignment. Here, the frame classification based APTAI model performed better, especially in terms of phoneme-related metrics. However, f -APTAI, based on forced alignment, has potentially more room for improvement in future work. An example of this, applicable to both models and requiring new pre-training, is changing the output frame rate of `wav2vec2` to 10 ms instead of 20 ms by changing the stride of the feature extractor, to improve alignment performance [23] and enable 100 Hz TV regression.

6. Acknowledgements

We gratefully acknowledge funding for this study by Friedrich-Alexander-University Erlangen-Nuermberg, Medical Valley e.V. and Siemens Healthineers AG within the framework of d.hip campus.

7. References

- [1] C. P. Browman and L. Goldstein, "Gestural specification using dynamically-defined articulatory structures," *Journal of Phonetics*, vol. 18, no. 3, pp. 299–320, 1990.
- [2] A. Ji, "Speaker independent acoustic-to-articulatory inversion," Ph.D. dissertation, Marquette University, 2014.
- [3] R. S. McGowan, "Recovering articulatory movement from formant frequency trajectories using task dynamics and a genetic algorithm: Preliminary model tests," *Speech Communication*, vol. 14, no. 1, pp. 19–48, 1994.
- [4] J. Chartier, G. K. Anumanchipalli, K. Johnson *et al.*, "Encoding of articulatory kinematic trajectories in human speech sensorimotor cortex," *Neuron*, vol. 98, no. 5, pp. 1042–1054, 2018.
- [5] P. Wu, L.-W. Chen, C. J. Cho, S. Watanabe, L. Goldstein, A. W. Black, and G. K. Anumanchipalli, "Speaker-independent acoustic-to-articulatory speech inversion," 2023.
- [6] J. Wang, J. Liu, L. Zhao, S. Wang, R. Yu, and L. Liu, "Acoustic-to-articulatory inversion based on speech decomposition and auxiliary feature," in *ICASSP 2022-2022*. IEEE, 2022, pp. 4808–4812.
- [7] Y. M. Siriwardena, G. Sivaraman, and C. Espy-Wilson, "Acoustic-to-articulatory speech inversion with multi-task learning," *arXiv preprint arXiv:2205.13755*, 2022.
- [8] G. Beguš, A. Zhou, P. Wu, and G. K. Anumanchipalli, "Articulation gan: Unsupervised modeling of articulatory learning," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [9] S. K. Maharana, A. Illa *et al.*, "Acoustic-to-articulatory inversion for dysarthric speech by using cross-corpus acoustic-articulatory data," in *ICASSP 2021*. IEEE, pp. 6458–6462.
- [10] N. R. Benway, Y. M. Siriwardena *et al.*, "Acoustic-to-Articulatory Speech Inversion Features for Mispronunciation Detection of /t/ in Child Speech Sound Disorders," in *Proc. INTERSPEECH 2023*, pp. 4568–4572.
- [11] C. Haldin *et al.*, "Speech rehabilitation in post-stroke aphasia using visual illustration of speech articulators: A case report study," *Clinical Linguistics & Phonetics*, vol. 35, no. 3, pp. 253–276, 2021.
- [12] T. Sweeney, F. Hegarty *et al.*, "Randomized controlled trial comparing parent led therapist supervised articulation therapy (plat) with routine intervention for children with speech disorders associated with cleft palate," *International Journal of Language & Communication Disorders*, vol. 55, no. 5, pp. 639–660, 2020.
- [13] Y. M. Siriwardena and C. Espy-Wilson, "The secret source: Incorporating source features to improve acoustic-to-articulatory speech inversion," in *ICASSP 2023*. IEEE, pp. 1–5.
- [14] N. Seneviratne, G. Sivaraman, and C. Espy-Wilson, "Multi-Corpus Acoustic-to-Articulatory Speech Inversion," in *Proc. Interspeech 2019*, 2019, pp. 859–863.
- [15] G. Sivaraman, V. Mitra, H. Nam, M. Tiede, and C. Espy-Wilson, "Unsupervised speaker adaptation for speaker independent acoustic to articulatory speech inversion," *The Journal of the Acoustical Society of America*, vol. 146, no. 1, pp. 316–329, 2019.
- [16] A. S. Shahrehabaki, S. M. Siniscalchi *et al.*, "Sequence-to-sequence articulatory inversion through time convolution of sub-band frequency signals," *Proc. Interspeech 2020*.
- [17] T. Biasutto-Lervat and S. Ouni, "Phoneme-to-articulatory mapping using bidirectional gated rnn," in *Interspeech 2018*.
- [18] V. Ribeiro, K. Isaieva *et al.*, "Automatic generation of the complete vocal tract shape from the sequence of phonemes to be articulated," *Speech Communication*, vol. 141, pp. 1–13, 2022.
- [19] A. Singh, A. Illa, and P. K. Ghosh, "A comparative study of estimating articulatory movements from phoneme sequences and acoustic features," in *ICASSP 2020*, pp. 7334–7338.
- [20] S. Udupa, A. Roy, A. Singh, A. Illa, and P. K. Ghosh, "Estimating articulatory movements in speech production with transformer networks," *Proc. Interspeech 2021*.
- [21] Q. Li, C. Zhang, and P. C. Woodland, "Combining frame-synchronous and label-synchronous systems for speech recognition," *arXiv preprint arXiv:2107.00764*, 2021.
- [22] R. Badlani, A. Łańcucki, K. J. Shih, R. Valle, W. Ping, and B. Catanzaro, "One tts alignment to rule them all," in *ICASSP 2022*. IEEE, 2022, pp. 6092–6096.
- [23] J. Zhu, C. Zhang, and D. Jurgens, "Phone-to-audio alignment without text: A semi-supervised approach," in *ICASSP 2022*. IEEE, pp. 8167–8171.
- [24] K. J. Shih, R. Valle *et al.*, "Rad-tts: Parallel flow-based tts with robust alignment learning and diverse synthesis," in *ICML Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models*, 2021.
- [25] F. Kreuk, J. Keshet, and Y. Adi, "Self-Supervised Contrastive Learning for Unsupervised Phoneme Segmentation," in *Proc. Interspeech 2020*, pp. 3700–3704.
- [26] Y. Li, B. J. Wohlan, D.-S. Pham, K. Y. Chan, R. Ward, N. Hennessey, and T. Tan, "Improving text-independent forced alignment to support speech-language pathologists with phonetic transcription," *Sensors*, vol. 23, no. 24, p. 9650, 2023.
- [27] J. Lian, C. Feng, N. Farooqi, S. Li *et al.*, "Unconstrained dysfluency modeling for dysfluent speech transcription and detection," in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pp. 1–8.
- [28] 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.
- [29] M. Parrot, J. Millet, and E. Dunbar, "Independent and automatic evaluation of acoustic-to-articulatory inversion models," *Proc. Interspeech 2020*.
- [30] T. tom Dieck, P.-A. Pérez-Toro, T. Arias-Vergara, E. Nöth, and P. Klumpp, "Wav2vec behind the scenes: How end2end models learn phonetics," *Proc. Interspeech 2022*, pp. 5130–5134.
- [31] 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.
- [32] T. Kislir, U. Reichel, and F. Schiel, "Multilingual processing of speech via web services," *Computer Speech & Language*, vol. 45, pp. 326–347, 2017.
- [33] P. Klumpp *et al.*, "Common phone: A multilingual dataset for robust acoustic modelling," *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pp. 763–768, 2022.
- [34] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer *et al.*, "Common voice: A massively-multilingual speech corpus," *arXiv preprint arXiv:1912.06670*, 2019.
- [35] J. S. Garofolo, "Timit acoustic phonetic continuous speech corpus," *Linguistic Data Consortium*, 1993, 1993.
- [36] M. Tiede, C. Y. Espy-Wilson *et al.*, "Quantifying kinematic aspects of reduction in a contrasting rate production task," *The Journal of the Acoustical Society of America*, vol. 141, no. 5_Supplement, pp. 3580–3580, 2017.