



Exploration strategies for articulatory synthesis of complex syllable onsets

Daniel R. van Niekerc¹, Anqi Xu¹, Branislav Gerazov², Paul K. Krug³, Peter Birkholz³, Yi Xu¹

¹Department of Speech, Hearing and Phonetic Sciences, University College London, UK

²Faculty of Electrical Engineering and Information Technologies, CMUS, Skopje, RN Macedonia

³Institute of Acoustics and Speech Communication, Technische Universität Dresden, Germany

d.vniekerk@ucl.ac.uk

Abstract

High-quality articulatory speech synthesis has many potential applications in speech science and technology. However, developing appropriate mappings from linguistic specification to articulatory gestures is difficult and time consuming. In this paper we construct an optimisation-based framework as a first step towards learning these mappings without manual intervention. We demonstrate the production of syllables with complex onsets and discuss the quality of the articulatory gestures with reference to coarticulation.

Index Terms: articulatory phonetics, articulatory speech synthesis, coarticulation, consonant clusters.

1. Introduction

High-quality articulatory speech synthesis provides compelling possibilities for studying articulatory phonetics and constructing low-resource speech technologies. However, many of these scenarios require a mapping from linguistic specification, e.g. a sequence of phonetic symbols to articulatory gestures. To date, such mappings have been developed manually by developers or users of the particular synthesizer. *VocalTractLab* [1, 2], a state-of-the-art synthesizer that simulates a vocal tract based on magnetic resonance imaging (MRI) data, is a system capable of producing natural sounding speech [3]. However, the increase in realism comes with an increase in time and expertise required to develop gestures and the increase in speech quality reflects subtle articulatory choices which makes it difficult to develop language or dialect-independent gestures manually.¹ The result is that the set of available gestures will be, at best, incomplete or in the worst case only appropriate for a specific language.

This problem can be resolved by a procedure that learns the articulatory gestures to produce linguistically relevant utterances automatically. To be practical, the process should function without detailed articulatory phonetic information such as aligned MRI data or expert knowledge. This resembles the task of *spoken language acquisition*, in general, and *early vocal learning* in particular [4]. In this paper we implement one of the central processes for autonomous learning of speech production, namely articulatory exploration or “babbling”, as an auditory optimisation task [5].

Using this simulation, we show that it is possible to discover the articulation of syllables with complex onsets by using a perceptual encoder trained on a general speech recognition corpus to provide auditory objectives. Furthermore we demonstrate the relative success of different exploration strategies and examine the nature of articulatory solutions in terms of coarticulation between consonant and vowel gestures.

¹Furthermore, gestural mappings will need to be revisited each time a new speaker model is introduced.

This is a significant development towards an autonomous process for constructing linguistic-to-articulatory mappings for any language or dialect and provides a framework for investigating theoretical questions in articulatory phonetics and speech perception in vocal learning.

2. Approach

As in previous work [6] we formulate articulatory exploration as an optimisation task using *VocalTractLab* (VTL) to produce candidate utterances. However, in this work, novel mechanisms including well-motivated somatosensory specifications and a language-oriented auditory perceptual mapping form part of the objective function. The process of discovering linguistically relevant gestural mappings is illustrated in Figure 1 and is briefly motivated in the following subsections.

2.1. Articulatory exploration

Babbling during early vocal learning has often been simulated as a *goal-directed or imitative* process, usually involving a set of auditory objectives [7–11]. This type of exploration is also considered central to finding appropriate inverse models during sensorimotor learning in general [12, 13].

We implement goal-directed articulatory exploration as the process of minimising auditory and somatosensory losses to discover a linguistically relevant utterance. The central block in Figure 1 is the global optimisation task

$$\mathbf{u}^* = \operatorname{argmin}_{\mathbf{u} \in \mathcal{U}_\theta} L(\mathbf{u}, \mathbf{q}, Q, \theta) \quad (1)$$

of finding the articulatory gestures \mathbf{u}^* that minimise the loss function L , with \mathbf{q} the combined auditory/articulatory goal, Q the auditory perceptual mapping described in Section 2.2 and θ the speaker vocal tract model. In this paper, we use the Tree-structured Parzen Estimator approach [14] as the algorithm to drive the articulatory sampling.

2.2. Auditory perceptual objectives

Most vocal learning simulations assume that auditory objectives are derived from the speech signal alone, i.e. acoustic imitation. This approach suffers from two issues: the *speaker normalisation* and *correspondence* problems. The speaker normalisation problem refers to the difficulty of finding linguistically correct utterances when comparing speech produced by different speakers; e.g. it is well known that formant frequencies vary systematically with speakers’ vocal tract length and that this may affect speech recognition performance [15]. The correspondence problem is one of associating articulatory gestures obtained for an acoustic reference to linguistic contexts [11, 16, 17]. We have argued elsewhere, based on the

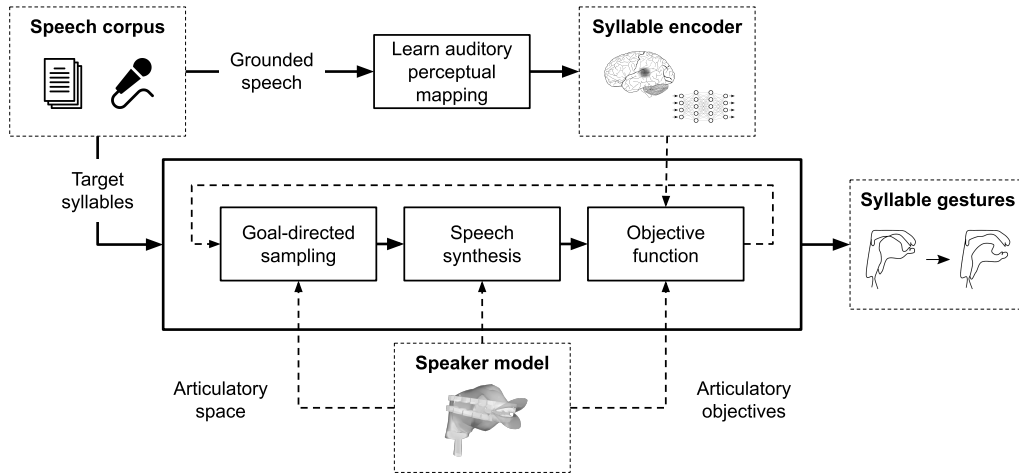


Figure 1: A process for discovery of linguistically relevant articulatory gestures. The central exploration task is goal-directed and relies on articulatory sampling, speech production, and auditory and articulatory objectives. Required input resources are a multi-speaker speech corpus and the vocal tract model.

well-known finding that language-oriented speech perception precedes the onset of *canonical babbling* in infants [18], that these problems can be addressed by a language-oriented auditory perceptual mapping derived from linguistically grounded multi-speaker speech stimuli [19].

2.3. Articulatory objectives

Articulatory objectives represent *explicit objectives* that originate from non-auditory signals. In humans it is known that sighted individuals may benefit from visual information [20] and speakers may track the implementation of these objectives through somatosensory feedback [21]. However, since we rely on a general optimisation algorithm with uniform priors instead of a physiologically motivated approach (as for example in [22, 23]) and the speaker model does not incorporate inertial and other relevant process measurements, articulatory objectives may also serve as an *implicit mechanism* that regularises the solution space, i.e. resulting in more prototypical articulatory gestures.

In this paper we employ one set of *somatosensory objectives* that could be derived from visual information: plosive consonants at the start of the syllable should form an oral closure and the vowel is associated with an open vocal tract. We also experiment with a *regularisation objective* to induce intra-syllable coarticulation.

2.4. Speech production

To produce articulatory trajectories, we use the target-approximation model (TAM) [24] which has been adopted in *VTL* to realise utterances represented by *articulatory targets* [25]. The resulting parameterisation of the articulatory dynamics combined with simplifying assumptions of synchronisation [26] has enabled the reliable discovery of simple CV syllables using derivative-free optimisation or even random sampling [6, 27]. While previous works implemented coarticulation [28, 29] by explicitly parameter tying [6, 27], this work tests the hypothesis by including coarticulation as an articulatory objective, which further reduces the explicit knowledge required in the process, as described in the next section.

Table 1: Upper vocal tract parameters in *VTL*.

Parameter	Description
HX, HY	Horiz. and vert. hyoid positions
JX, JA	Horiz. jaw position and jaw angle
LP, LD	Lip protrusion and vert. lip distance
TTX, TTY	Horiz. and vert. tongue tip positions
TBX, TBY	Horiz. and vert. tongue blade positions
TCX, TCY	Horiz. and vert. tongue body centre positions
TRX, TRY	Horiz. and vert. tongue root positions
VS, VO	Velum shape and opening
TS1 – TS3	Tongue side elevation from the anterior to the posterior part of the tongue

3. Experimental setup

3.1. Implementation

The first step towards implementing the articulatory exploration process (Figure 1) to find CCV syllables is to construct the auditory perceptual mapping that produces syllable embeddings or percepts. For this purpose, we used the *LibriSpeech* speech recognition corpus [30] representing linguistically grounded speech. Vowel (V), consonant-vowel (CV) and CCV syllable onsets were extracted from the clean training set and used to train a recurrent neural network that encodes the audio to a low-dimensional space related to the linguistic context. The Mel-spectrogram was used as input and the output vector was a concatenation of one-hot encoded phonetic identities defined in the ARPABET phoneset (used in the CMU pronunciation dictionary [31]) which is appropriate for the American English speech data. The resulting vector $[q_{c1}, q_{c2}, q_v]$ representing a V, CV or CCV syllable onset is 64-dimensional – two sub-vectors encoding 24 consonants (including absence) and one representing 16 vowels. Evaluating this encoder on the Librispeech test set by converting the output to a categorical form results in a recognition rate of 73%.

For speech synthesis, *VTL*² was used to realise articulatory targets with the “JD2” male speaker and geometric glottis

²Version 2.3 available at <https://www.vocaltractlab.de>

Table 2: Identification rates (%) for samples by vowel, onset, and complete syllable for each experimental condition (5 trials \times 150 syllable types = 750 samples per condition).

	Without coarticulation objective				With coarticulation objective			
	C_1C_2V	$V.C_1C_2$	$V.C_1.C_2$	$V.C_2.C_1$	C_1C_2V	$V.C_1C_2$	$V.C_1.C_2$	$V.C_2.C_1$
syllable	83.60	83.87	82.93	<u>69.47</u>	79.97	85.07	81.15	<u>70.53</u>
vowel	91.33	93.07	92.27	<u>90.53</u>	<u>88.79</u>	94.00	92.91	92.93
onset	91.07	90.53	90.27	<u>77.47</u>	89.99	90.93	87.83	<u>76.80</u>

model [32]. Since we were focused on investigating the upper vocal tract parameters, the glottal parameters were kept constant at the appropriate preset values for the particular segment (e.g. “modal voice” for the vowel), with the exception of the *chink area* and *relative amplitude* which were free to be optimised to allow control of the voice onset time. All of the upper vocal tract parameters (Table 1) were free for optimisation, except the *velum opening* (VO) which was kept closed and the *tongue root* (TRX, TRY) parameters which were derived from the tongue body values [33]. Timing in the target-approximation trajectories was controlled by two free parameters, one time constant each for the glottal and upper vocal tract parameters.

The somatosensory objectives relied on the same VTL configuration to provide proprioceptive or tactile feedback by means of the *tube areas function*. Two simple objectives were defined and represented with values in the range [0, 1]. (1) The *vocal tract closure* objective value is 0 when a vocal tract closure is required, e.g. for plosive consonants, and 1 when a minimum opening is required, e.g. for vowels. (2) The *lip closure* objective value is 0 when a closure is formed by the lips and 1 when the lips are open. This is motivated by visual information and was only applied to the first consonant (C_1) depending on whether it is a bilabial or other type of plosive. No somatosensory objectives were applied to the intermediate consonant (C_2).

A single regularisation objective was implemented by quantifying the coarticulation between any two articulatory targets as the normalised L_1 distance between the range-normalised upper vocal tract vectors $\tilde{\mathbf{u}}$:

$$N^{-1} \|\tilde{\mathbf{u}}_1 - \tilde{\mathbf{u}}_2\|_1, \text{ where } \tilde{\mathbf{u}}_{1_i} \text{ and } \tilde{\mathbf{u}}_{2_i} \in [0, 1] \quad (2)$$

with N the number of upper vocal tract dimensions. In Section 4 we specifically compare systems with and without this *coarticulation objective* between each consonant and the vowel. Finally, all the relevant articulatory objectives were concatenated into a single vector \mathbf{q}_a and combined with the auditory percept to form the optimisation goal $[\mathbf{q}_a, \mathbf{q}_{c1}, \mathbf{q}_{c2}, \mathbf{q}_v]$.

To implement the optimisation algorithm [14] we used the *hyperopt*³ software package. The articulatory space was defined by the speaker model and initially sampled uniformly. The loss function was defined as the weighted sum of the Euclidean distances calculated for the individual sub-vectors, with auditory components having a weight ratio of 2:1 to the articulatory component. For each distinct syllable an ideal objective vector was constructed based on its phonetic constituents. The optimisation algorithm samples articulatory targets which are synthesised by VTL and each sample is evaluated by the auditory perceptual mapping and VTL tube areas function to determine the resulting vector and associated loss. To improve the computational efficiency of the process, the synthesis of speech and auditory evaluation is only performed when the somatosensory objectives are

satisfied. In the case of failure to achieve these objectives, the loss function is set to an arbitrary large value proportional to the loss associated with \mathbf{q}_a .

3.2. Evaluation

The proposed framework was evaluated through an experiment designed to investigate the following aspects:

1. *Exploration strategies* for complex syllable onsets: Is it necessary to optimise certain segment targets jointly or can this be done independently and in what sequence?
2. *Coarticulation*: Can we make use of the regularisation objective defined above to reproduce natural observations associated with intra-syllable coarticulation [29]?
3. *Sufficiency*: What is the relative success rate of the process for the range of CCV syllable types occurring in American English and what are the implications for future work?

Since the process is non-deterministic, dependent on random initial exploration, we estimated the success rate using independent repeated trials. To allow for comparison of different syllable types (aspect 3) we set up 5 trials for each of the 150 valid combinations⁴ of the following sets of segments: $C_1 \in /b,d,f,g,k,p,s,t,\theta/$, $C_2 \in /k,l,p,r,t,w/$, and $V \in /a,\ae,\a,\e,\o,i,i:,u,u:/$; a total of 750 independent trials for each experimental setup. Each trial was allocated 5000 iterations leading to a synthesised utterance, i.e. excluding articulatory targets that do not satisfy the basic somatosensory objectives. Four different exploration strategies were investigated (aspect 1):

1. *Single-pass, joint* (C_1C_2V): Find all segment targets jointly and select the best sample after 5000 evaluations.
2. *Two-pass, vowel then onset* ($V.C_1C_2$): Find the vowel targets by producing vowel-only utterances in a first pass, then find the onset consonants jointly by producing CCV utterances using the best vowel targets from the first pass.⁵ The number of iterations are allocated to the two passes in the ratio 1:4.
3. *Three-pass, with C_1 first* ($V.C_1.C_2$): Find the best targets for each segment by producing V , C_1V and C_1C_2V utterances in respective passes (iterations are allocated in the ratio 1:2:2).
4. *Three-pass, with C_2 first* ($V.C_2.C_1$): As in the previous configuration, but with the intermediate consonant explored first.

⁴Determined by the existence of entries in the CMU dictionary and the LibriSpeech corpus.

⁵To allow for varying coordination requirements in different contexts, the glottal parameters and time constants are never fixed but re-optimised during each pass.

³<https://github.com/hyperopt/hyperopt> (v0.2.5)




Figure 2: Mean articulatory distances (Eq. 2) between C_1 and V for the $V.C_1C_2$ condition with/without the coarticulation objective.

Each of the above strategies were implemented with and without the coarticulation objective (aspect 2) and the best outcome from each trial was evaluated by the syllable encoder. This was done by mapping perceptual representations to symbols using the argmax operation on each sub-vector $[q_{c1}, q_{c2}, q_v]$ and calculating the identification rate.

4. Results

The overall results are summarised in Table 2 in terms of identification rates, from which we note the following: (1) There is no significant difference⁶ in the auditory success rate when comparing conditions with and without the coarticulation objective. This means that adding the additional articulatory objective has no negative effect on the goal of optimising for auditory perception. (2) In both cases exploring C_1 last on the basis of a pre-optimised C_2V utterance leads to significantly worse results (underlined). (3) When applying the coarticulation objective, jointly optimising the vowel and consonants results in significantly worse outcomes for the vowel (C_1C_2V underlined) and optimising the consonants jointly after the vowel results in a significantly better outcome for syllables ($V.C_1C_2$ in bold).

By repeating this analysis over onsets and vowels we find that the outcomes with and without coarticulation in terms of auditory identification rate are similar in all contexts. In general, all contexts, except /d/, dw/, t/, tw/ for onsets and /u/, i:/ for vowels, have identification rates in excess of 80%, indicating that those cases are particularly difficult to discover.

To confirm that the coarticulation objective has the intended effect, articulatory distances are visualised in Figure 2. We see that: (1) There are significant reductions in the distance between C_1 and V of some articulatory parameters in each case. (2) Some expected patterns of articulatory overlap emerge, e.g. the bilabial targets have lower distances to the vowel in many dimensions except for the lip distance (LD).

5. Discussion

Since the optimisation process is partially dependent on the auditory perceptual mapping we are interested in (Table 2), the results should be interpreted carefully. The absolute identifica-

⁶We used Welch’s unequal variances t-test with a significance level of 5% throughout.

tion rates are not directly comparable with the recognition rates obtained on natural speech. Instead, we make use of the results in the following ways: (1) To compare the relative success of different exploration strategies in terms of the auditory objectives. (2) To indicate problematic contexts that require further work. The difficult cases listed in Section 4 could be addressed in future by additional somatosensory objectives or more ecologically plausible articulatory sampling in the case of the onsets and better modelling of duration in the case of the vowels.

The current optimisation-based simulation of babbling is a first step in learning articulation without manual intervention and will be supplemented by a gradient-based learning process towards fluent articulation in future. We invite the interested reader to listen to the samples available at https://github.com/danielshaps/evoclearn_optccv_2022 to get a qualitative sense for the extent to which these correspond to actual babbling utterances.

6. Conclusion

We have presented a flexible simulation of babbling that consists of auditory perceptual, somatosensory and regularisation objectives and demonstrated that it can discover the articulation of syllables with complex onsets. This framework was used to compare different exploration strategies leading to an effective process where the vowel targets are found independently in a first pass and the consonants in the onset are jointly optimised using the vowel as an “anchor”. With this two-pass procedure it is possible to apply the coarticulation objective (Eq. 2) without negatively affecting the outcomes in terms of the auditory perceptual goals. This means that the framework and analysis presented in Figure 2 can be used to discover the relative (in)dependence of articulatory dimensions in different contexts automatically – parameter tying was done manually in previous work [6,27]. Furthermore, this form of regularisation could prove useful if the utterances generated here are used as a basis for learning forward and inverse models of articulation [12].

7. Acknowledgements

This work has been funded by the Leverhulme Trust Research Project Grant RPG-2019-241: “High quality simulation of early vocal learning”.

8. References

- [1] P. Birkholz, *3D-Artikulatorische Sprachsynthese*. Berlin: Logos Verlag, 2005.
- [2] P. Birkholz, "Modeling Consonant-Vowel Coarticulation for Articulatory Speech Synthesis," *PLoS ONE*, vol. 8, no. 4, Apr. 2013.
- [3] P. K. Krug, S. Stone, and P. Birkholz, "Intelligibility and naturalness of articulatory synthesis with VocalTractLab compared to established speech synthesis technologies," in *Proc. 11th ISCA Speech Synthesis Workshop (SSW 11)*, pp. 102–107.
- [4] P. W. Jusczyk, *The Discovery of Spoken Language*. Cambridge, MA, USA: MIT Press, Jan. 1997.
- [5] D. K. Oller and R. E. Eilers, "The Role of Audition in Infant Babbling," *Child Development*, vol. 59, no. 2, pp. 441–449, 1988.
- [6] D. R. van Niekerk, A. Xu, B. Gerazov, P. K. Krug, P. Birkholz, and Y. Xu, "Finding intelligible consonant-vowel sounds using high-quality articulatory synthesis," in *Proc. Interspeech 2020*, Shanghai, China, Oct. 2020, pp. 4457–4461.
- [7] G. Bailly, "Learning to speak. Sensori-motor control of speech movements," *Speech Communication*, vol. 22, no. 2, pp. 251–267, Aug. 1997.
- [8] O. Räsänen, "Computational modeling of phonetic and lexical learning in early language acquisition: Existing models and future directions," *Speech Communication*, vol. 54, no. 9, pp. 975–997, Nov. 2012.
- [9] H. Rasilo and O. Räsänen, "An online model for vowel imitation learning," *Speech Communication*, vol. 86, pp. 1–23, Feb. 2017.
- [10] S. Pagliarini, A. Leblois, and X. Hinaut, "Vocal Imitation in Sensorimotor Learning Models: a Comparative Review," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 13, no. 2, pp. 326–342, Jun. 2021.
- [11] A. Philippsen, "Goal-Directed Exploration for Learning Vowels and Syllables: A Computational Model of Speech Acquisition," *KI - Künstliche Intelligenz*, Jan. 2021.
- [12] M. I. Jordan and D. E. Rumelhart, "Forward models: Supervised learning with a distal teacher," *Cognitive Science*, vol. 16, no. 3, pp. 307–354, Jul. 1992.
- [13] M. Rolf, J. J. Steil, and M. Gienger, "Goal Babbling Permits Direct Learning of Inverse Kinematics," *IEEE Transactions on Autonomous Mental Development*, vol. 2, no. 3, pp. 216–229, Sep. 2010.
- [14] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, "Algorithms for Hyper-Parameter Optimization," in *Proceedings of the 24th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, Dec. 2011, pp. 2546–2554.
- [15] P. Zhan and A. Waibel, "Vocal tract length normalization for large vocabulary continuous speech recognition," Carnegie Mellon University, Language Technologies Institute, Pittsburgh, PA, USA, Technical Report CMU-LTI-97-150, 1997.
- [16] C. L. Nehaniv and K. Dautenhahn, "The correspondence problem," in *Imitation in animals and artifacts*. Cambridge, MA, USA: MIT Press, Jan. 2002, pp. 41–61.
- [17] M. Brass and C. Heyes, "Imitation: is cognitive neuroscience solving the correspondence problem?" *Trends in Cognitive Sciences*, vol. 9, no. 10, pp. 489–495, Oct. 2005.
- [18] P. K. Kuhl, "Early language acquisition: cracking the speech code," *Nature Reviews Neuroscience*, vol. 5, no. 11, pp. 831–843, Nov. 2004.
- [19] D. R. Van Niekerk, A. Xu, B. Gerazov, P. K. Krug, P. Birkholz, and Y. Xu, "Vocal learning of spoken language: Beyond imitation," *Submitted to Speech Communication*, 2022.
- [20] M. Murakami, B. Kröger, P. Birkholz, and J. Triesch, "Seeing [u] aids vocal learning: Babbling and imitation of vowels using a 3D vocal tract model, reinforcement learning, and reservoir computing," in *International Conference on Development and Learning and on Epigenetic Robotics (ICDL-EpiRob)*, Providence, Rhode Island, USA, Aug. 2015, pp. 208–213.
- [21] S. M. Nasir and D. J. Ostry, "Somatosensory Precision in Speech Production," *Current Biology*, vol. 16, no. 19, pp. 1918–1923, Oct. 2006.
- [22] J. E. Serkhane, J. L. Schwartz, L. J. Boë, B. L. Davis, and C. L. Matyear, "Infants' vocalizations analyzed with an articulatory model: A preliminary report," *Journal of Phonetics*, vol. 35, no. 3, pp. 321–340, Jul. 2007.
- [23] H. Nam, L. M. Goldstein, S. Giulivi, A. G. Levitt, and D. H. Whalen, "Computational simulation of CV combination preferences in babbling," *Journal of Phonetics*, vol. 41, no. 2, pp. 63–77, Mar. 2013.
- [24] Y. Xu and Q. E. Wang, "Pitch targets and their realization: Evidence from Mandarin Chinese," *Speech Communication*, vol. 33, no. 4, pp. 319–337, Mar. 2001.
- [25] P. Birkholz, "Control of an Articulatory Speech Synthesizer Based on Dynamic Approximation of Spatial Articulatory Targets," in *Proc. Interspeech*, Antwerp, Belgium, Aug. 2007, pp. 2865–2868.
- [26] P. Birkholz, B. J. Kröger, and C. Neuschaefer-Rube, "Model-Based Reproduction of Articulatory Trajectories for Consonant-Vowel Sequences," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 5, pp. 1422–1433, Jul. 2011.
- [27] A. Xu, P. Birkholz, and Y. Xu, "Coarticulation as synchronized dimension-specific sequential target approximation: An articulatory synthesis simulation," in *Proceedings of the International Congress of Phonetic Sciences (ICPhS)*, Melbourne, Australia, Aug. 2019, pp. 205–209.
- [28] Y. Xu, "Syllable is a synchronization mechanism that makes human speech possible," *PsyArXiv*, Mar. 2020. [Online]. Available: <https://osf.io/9v4hr>
- [29] Z. Liu, Y. Xu, and F.-F. Hsieh, "Coarticulation as synchronised CV co-onset - Parallel evidence from articulation and acoustics," *Journal of Phonetics*, in-press.
- [30] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2015, pp. 5206–5210.
- [31] Carnegie Mellon University. (2000) The CMU pronunciation dictionary. [Online]. Available: <http://www.speech.cs.cmu.edu/>
- [32] P. Birkholz, S. Drechsel, and S. Stone, "Perceptual optimization of an enhanced geometric vocal fold model for articulatory speech synthesis," in *Proc. Interspeech*, Graz, Austria, Sep. 2019, pp. 3765–3769.
- [33] P. K. Krug, P. Birkholz, B. Gerazov, D. R. van Niekerk, A. Xu, and Y. Xu, "Efficient Exploration of Articulatory Dimensions," *Studientexte zur Sprachkommunikation: Elektronische Sprachsignalverarbeitung 2022*, pp. 51–58, 2022.