

Automatic Measurement of Pre-aspiration

Yaniv Sheena¹, Míša Hejna², Yossi Adi¹, Joseph Keshet¹

¹Department of Computer Science, Bar-Ilan University, Ramat-Gan, Israel

²Department of English, Aarhus University, Denmark

yaniv.sheena@live.biu.ac.il, misa.hejna@cc.au.dk, jkeshet@cs.biu.ac.il

Abstract

Pre-aspiration is defined as the period of glottal friction occurring in sequences of vocalic/consonantal sonorants and phonetically voiceless obstruents. We propose two machine learning methods for automatic measurement of pre-aspiration duration: a feedforward neural network, which works at the frame level; and a structured prediction model, which relies on manually designed feature functions, and works at the segment level. The input for both algorithms is a speech signal of an arbitrary length containing a single obstruent, and the output is a pair of times which constitutes the pre-aspiration boundaries. We train both models on a set of manually annotated examples. Results suggest that the structured model is superior to the frame-based model as it yields higher accuracy in predicting the boundaries and generalizes to new speakers and new languages. Finally, we demonstrate the applicability of our structured prediction algorithm by replicating linguistic analysis of pre-aspiration in Aberystwyth English with high correlation.

Index Terms: pre-aspiration, feedforward neural network, structured prediction, laboratory phonology

1. Introduction

Pre-aspiration is a period of (primarily) glottal friction, which is found in the sequences of sonorants and phonetically voiceless obstruents prior to the release of the consonant, e.g., in Welsh English *kit* /k^hi^ht^s/, *miss* /mi^hs/, *milk* /mi^hlk^h/, etc. Although pre-aspiration was previously considered rare [1,2], with the advances in recording technology it has been recently reported to occur in increasingly more languages and varieties thereof [3–15]. Recent studies of pre-aspiration have revealed that the phenomenon is very individual [1, 7, 16], sensitive to sex/gender [17], and can be affected by age [5, 7] and the first language of the speaker [17].

Most of the work on pre-aspiration has been based on subjective, labor-intensive manual annotation. The vast majority of researchers have coded pre-aspiration manually by identifying its boundaries, usually based on a manual segmentation of the signal into segmental and subsegmental intervals [1, 4, 7, 9, 10, 12, 13, 17–21].

In this paper we tackle the problem of automatic measurement of pre-aspiration duration. We propose two algorithms for this purpose. The first one is a feedforward neural network, which works at the frame-level, and the second algorithm is based on structured prediction techniques which rely on highly tuned feature maps which were specifically designed for the task, and which works at the segment level.

As far as we know this is the first attempt to automatically detect pre-aspiration. This work is based on our ongoing work on designing and developing machine learning algorithms for automatically measuring with high precision, phonetic properties of speech at the level of human inter-transcriber reliabil-




Figure 1: Annotation example of pre-aspiration. The signal consists of a vowel followed by a plosive of the word “cook”.

ity [22–26]. Our methods rely on several advances over existing computational systems: novel representations of the speech signal and new structured prediction and deep learning algorithms. These automatic methods allow for low-cost replication of phonetic studies and expand the range of empirical and theoretical issues we can address.

The code is available to be used by the research community. It can be downloaded from <https://github.com/MLSpeech/AutoPreaspiration>.

2. Problem Setting

In the context of a typical phonological study, given a segment of an acoustic signal of an arbitrary length, the goal of an automatic pre-aspiration measurement algorithm is to predict the onset and offset times of pre-aspiration as accurate as possible. In this work we assume that the acoustic signal contains exactly one coded obstruent within a word and the signal starts before the expected pre-aspiration, i.e., during the previous phoneme.

We turn to describing the problem formally. Throughout the paper, scalars are denoted using lower case Latin letters, e.g., x , and vectors using bold face letters, e.g., \mathbf{x} . A sequence of elements is denoted with a bar \bar{x} and its length is written as $|\bar{x}|$.

The acoustic input is represented as a sequence of feature vectors, $\bar{\mathbf{x}} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$, where each vector, \mathbf{x}_t , is D -dimensional vector which represents the acoustic content of the t -frame ($1 \leq t \leq T$). The domain of the feature vectors is denoted as $\mathcal{X} \subset \mathbb{R}^D$. The acoustic input is of an arbitrary length, hence the number of frames, T , is not fixed. We denote by \mathcal{X}^* the set of all finite-length sequences over \mathcal{X} .

Each acoustic input is associated with a timing pair: the *onset* of the pre-aspiration, $t_s \in \mathcal{T}$, and the *offset* of the pre-aspiration, $t_e \in \mathcal{T}$, where $\mathcal{T} = \{1, \dots, T\}$. To sum up, given the speech interval $\bar{\mathbf{x}}$, our goal is to learn a function f from the domain of all acoustic inputs \mathcal{X}^* to the domain of all possible onset offset pairs \mathcal{T}^2 . Our notation is depicted in Figure 1.

3. Learning Apparatus

In this section, we describe how we learn the prediction function from a training set of examples. We denote the training set of

m examples by $S = \{(\bar{\mathbf{x}}^i, t_s^i, t_e^i)\}_{i=1}^m$, where the i -th example is composed of a sequence of acoustic features $\bar{\mathbf{x}}^i$ labeled with an onset and offset pair, (t_s^i, t_e^i) .

Let us denote the predicted onset and offset pair by (\hat{t}_s, \hat{t}_e) , namely $(\hat{t}_s, \hat{t}_e) = f_\theta(\bar{\mathbf{x}})$, where θ denotes the set of parameters of the prediction function f . In order to assess the quality of the prediction we use a *task loss* function, $\gamma((t_e, t_s), (\hat{t}_e, \hat{t}_s))$, which returns a real positive number that measures how the prediction pair (\hat{t}_s, \hat{t}_e) is close to the manually annotated pair (t_e, t_s) . Our goal in learning is to find the set of parameters θ so as to minimize the expected task loss.

We start by describing the acoustic features, and then two different machine learning algorithms that learn the parameters of function f . The first algorithm works at the frame level and hence cannot aim at minimizing a global task loss function, while the second algorithm is aimed directly at minimizing the expected task loss. These differences are reflected in the accuracy level of the different algorithms.

3.1. Acoustic features

For both machine learning models, we extract the same set of features. Consider the acoustic input $\bar{\mathbf{x}} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$ consisting of T frames, where each acoustic feature vector \mathbf{x}_t consists of D features. Similar to [22], we extract eight ($D=8$) acoustic features from the speech signal every 1 ms. The first four features refer to the total spectral energy (E_{total}), energy between 50-1000 Hz (E_{low}), energy above 3000 Hz (E_{high}), and Wiener entropy (H_{wiener}) — all are based on the short-time Fourier transform (STFT) taken every 1 ms with a 5 ms Hamming window. The fifth feature, P_{max} , is the maximum of the power spectrum calculated in a region from 6 ms before to 18 ms after the frame center. The sixth feature (R_i) is the pitch of the signal using a real-time pitch detector [27]. The seventh feature is the 0/1 output of a voicing detector based on the RAPT pitch tracker [28], smoothed with a 5 ms Hamming window (V). The last feature is the number of zero crossings (ZC) in a 5 ms window around the frame center.

3.2. Frame-based model

Our first model works at the frame level. We train a binary classifier such that given an input speech frame \mathbf{x}_t predicts whether or not it is associated with a pre-aspiration event. In order to take advantage of the local temporal context of each time-frame, the input of this binary classifier is based on five concatenated feature vectors $(\mathbf{x}_{t-2}, \mathbf{x}_{t-1}, \mathbf{x}_t, \mathbf{x}_{t+1}, \mathbf{x}_{t+2})$ rather than a single frame. We use a feedforward neural network. We tried several network architectures, and chose the one that performed best on a validation set. The network consists of an input layer of 40 units (recall that we extract $D = 8$ feature per frame, and have 5 consecutive frames), a hidden layer of 40 units with ReLU activation function and a dropout rate of 0.3, and one output unit followed by a sigmoid. The network was trained to minimize the binary cross entropy loss using the gradient descent optimization algorithm.

At inference time, we use the trained model sequentially over the input frames to produce a sequence of predictions. Since this sequence can be noisy, we smooth it out using a moving average, followed by a binary mapping that uses a threshold for each such average. In the final step, we search the longest subsequence of frames that were predicted as associated with pre-aspiration, and output its boundaries as the onset and offset of the pre-aspiration.

3.3. Structured model

The second algorithm takes advantage of the input as a whole segment, hence we can introduce feature maps such as typical pre-aspiration duration, mean energy during the presumed pre-aspiration compared to the mean energy before or after the pre-aspiration. Similar to previous work in structured prediction [29, 30], the function f is constructed from a predefined set of N feature maps $\{\phi_i\}_{i=1}^N$, each of the form $\phi_i : \mathcal{X} \times \mathcal{T} \times \mathcal{T} \times \mathbb{R}^N \rightarrow \mathbb{R}^N$, and a weight vector $\mathbf{w} \in \mathbb{R}^N$. The function is a linear predictor of the following form

$$(\hat{t}_s, \hat{t}_e) = \arg \max_{(t_s, t_e)} \mathbf{w} \cdot \phi(\bar{\mathbf{x}}, t_s, t_e), \quad (1)$$

where we have used vector notation for the feature maps $\phi = (\phi_1, \dots, \phi_N)^T$. This vector-valued function is used to map the variable length of input speech along with a presumed onset-offset pair to an abstract vector space in \mathbb{R}^N (described in 3.4).

The algorithm presented here aims to find the weights \mathbf{w} that minimize the expected task loss function γ which measures the distance between predicted and manually coded labels. In a similar manner to [22], we define the task loss function as follows:

$$\gamma((t_e, t_s), (\hat{t}_e, \hat{t}_s)) = \max\{|\hat{t}_e - \hat{t}_s - (t_e - t_s)| - \epsilon, 0\}. \quad (2)$$

That is, only the differences between the predicted pre-aspiration and the manually labeled pre-aspiration that are greater than a threshold ϵ (in milliseconds), are penalized. This task loss function takes into account that manual measurements are not usually exact, and ϵ can be adjusted according to the level of measurement uncertainty in a dataset.

The weight vector \mathbf{w} is learned using an iterative algorithm based on the *Passive-Aggressive* family of algorithms for structured prediction [31]. Let \mathbf{w}_t be the weight vector after the t -th iteration, and let $\mathbf{w}_0 = \mathbf{0}$. At each iteration a single example $(\bar{\mathbf{x}}^i, t_s^i, t_e^i)$ is considered, and the current weight vector \mathbf{w}_t is updated by finding the solution to the following optimization problem:

$$\begin{aligned} \mathbf{w}_{t+1} &= \arg \min_{\mathbf{w}, \xi \geq 0} \frac{1}{2} \|\mathbf{w} - \mathbf{w}_t\|^2 + C\xi \\ \text{s.t. } &\mathbf{w} \cdot \phi(\bar{\mathbf{x}}^i, t_s^i, t_e^i) - \mathbf{w} \cdot \phi(\bar{\mathbf{x}}^i, \tilde{t}_s, \tilde{t}_e) \geq \gamma_i - \xi, \end{aligned} \quad (3)$$

where $\gamma^i = \gamma((t_s^i, t_e^i), (\tilde{t}_s, \tilde{t}_e))$, C serves as a trade-off parameter between loss and regularization minimization, and ξ is a non-negative slack variable.

The optimization problem tries to keep the new weight vector \mathbf{w} close to the previous weight vector \mathbf{w}_t while satisfying the constraint that the score of the manually annotated onset-offset pair $\mathbf{w} \cdot \phi(\bar{\mathbf{x}}^i, t_s^i, t_e^i)$ will be higher than a different set of onset-offset, $\mathbf{w} \cdot \phi(\bar{\mathbf{x}}^i, \tilde{t}_s, \tilde{t}_e)$, where the onset-offset pair $(\tilde{t}_s, \tilde{t}_e)$ in the constraint is the most violated pair [30], namely:

$$(\tilde{t}_s, \tilde{t}_e) = \arg \max_{(t_s, t_e)} \mathbf{w} \cdot \phi(\bar{\mathbf{x}}^i, t_s, t_e) + \gamma^i. \quad (4)$$

3.4. Feature maps for the structured model

The feature maps are built from local differences, cumulative mean and max over subsets features with respect to the pre-aspiration boundaries. They were chosen based on manual inspection of the values and the trends of the features in intervals

Table 1: Summary of the feature maps. The rows represent the type of calculations that should be performed. The columns are the eight acoustic features. F in row i and column j indicates that there is a feature map of type i for feature x_j ; Δ indicates that there are three feature maps of type i for the local difference of feature x_j , evaluated at $s = 5, 10, 15$ [22]. For example, the F in the last row, column E_{high} denotes a feature-map that gets the mean of the value of E_{high} over 50 frames starting from t_e .

Feature map type	E_{total}	E_{low}	E_{high}	H_{wiener}	P_{max}	R_t	V	ZC
Value at t_s	F, Δ	Δ	F, Δ	F, Δ	Δ	F, Δ	Δ	Δ
Value at t_e	F, Δ	F, Δ	F, Δ	F, Δ	F, Δ	F, Δ		
Mean & max over (t_s, t_e)					F, Δ_5, Δ_{10}			
Mean over (t_s, t_e) - mean over $(t_e, t_e + 50)$			F	F				F
Mean over $(t_e, t_e + 50)$			F	F		F	F	
Max over $(t_e, t_e + 50)$			F	F				

near t_s and t_e , and they reflect the knowledge about the problem of pre-aspiration measurement, which in our case is based on [7].

As an example for such a feature map we considered the observation that when pre-aspiration is immediately followed by a silent interval in the context of plosives, we expect that high frequencies (above 3000 Hz) will decrease compared to the interval of pre-aspiration, which has high energy presence in these frequencies. In order to express this observation, we define feature maps that compute the differences of the means of E_{high} and H_{wiener} over a post pre-aspiration interval $(t_e, t_e + 50)$ and the pre-aspiration interval (t_s, t_e) .

Another key observation is that t_s usually comes right after a formant structure which becomes less distinct and usually indicates that voicing is ending. Hence a set of feature maps are based on the local differences of 5, 10 and 15 milliseconds over the acoustic features P_{max} and V in order to allow our algorithm to capture these changes. A full specification of ϕ is shown in Table 1.

4. Data Set

The data is composed of 5,297 examples of 16 speakers of English from the town of Aberystwyth within mid Wales. All of the speakers were born and raised in the town and are L1 Welsh speakers. The speakers presented are 10 females and 6 males, and their age range spans from 22–91 to allow for preliminary generational comparisons. The data itself consisted of a list of words that the participants read; each word contained a sequence of a vowel and a post-tonic plosive. Each word was read once in isolation and twice in a frame sentence. For more details on the segmental and prosodic characteristics of the tokens related to aspects such as vowel height, place and manner of articulation of the plosive, stress, and foot-position see [7]. The recordings were obtained with an H4 Zoom Handy Recorder (sampled at 44.1kHz) in conjunction with the head-mounted AKG C520 Microphone, which was attached to the speaker’s head and ensured a constant distance from the mouth, irrespective of the speaker’s movements.

5. Experiments

We evaluated the performance of our algorithms using three main methods. First, we compared the predictions of the frame-based algorithm and of the structured algorithm to the manual annotations. Then we tested how the structured prediction algorithm, which was superior to the frame-based algorithm, generalizes on new set of speakers and languages. Finally, we replicated a linguistic study on the model predictions and compared it to the results obtained from the manually annotated data. All of the experiments in this section are based on models which

were trained using search windows of 50 ms before the labeled left boundary (t_s) and 60 ms after the right boundary (t_e), where we tried several window-sizes and the effects were insignificant.

5.1. Models performance

We evaluated both algorithms on the corpus described in the previous section using 5-fold cross-validation. For each fold, 15% of the data served as a validation set. We regularized both algorithms using early stopping. For the frame-based algorithm, we balanced the data by randomly dropping negative examples, as the amount of such examples was significantly greater than the amount of positive examples. For the structured model we used $C = 50$, which was chosen on a validation set, and ϵ in the task loss was set to 2 ms.

We compared the models by reporting the percentage of examples in the test set with automatic/manual differences which were less than a time tolerance, where we used tolerances of 5, 10, 15, and 20 ms. We also report the average error in the prediction of the onset and the offset. The results are given in Table 2. Results should read as follows: The percentage of correctly predicted pre-aspirations within a threshold of 2 ms was 43.3% for the frame-based algorithm and 56.2% for the structured algorithm. The mean difference between the predicted and manually-annotated onset was 4.4 ms for the frame-based algorithm and 2.6 ms for the structured algorithm.

Table 2: Results for the two algorithms. The first 4 columns are the percentage of accurately predicted pre-aspirations within a given threshold (thresholds in ms, values are percentages). The remaining columns are the mean distance between the boundaries of manual/automatic pre-aspirations in ms.

Algorithm	≤ 5	≤ 10	≤ 15	≤ 20	Δt_s	Δt_e
Frame-based	43.3	74.3	88.9	95.3	4.4	5.3
Structured	56.2	84.8	93.1	96.3	2.6	4.9

In Table 3 we report the mean and the standard deviation of the manually annotated pre-aspiration duration and the predicted pre-aspiration duration for both algorithms.

Table 3: Mean and standard deviation of pre-aspiration duration. Manually coded versus models’ predictions.

Type	Mean (ms)	Standard deviation (ms)
Manual	37.6	20.8
Structured	37.2	19.9
Frame-based	42.0	17.8

Results suggest that the structured algorithm is superior to the framed-based algorithm. This demonstrates the power of the structured algorithm, which uses a set of dedicated feature-maps each of which measures properties and trends of intervals surrounding the boundaries of the pre-aspiration, contrary to the frame-based method that is based on local windows. Henceforth, we conducted the rest of the experiments using only the structured algorithm.

5.2. Generalization

In order to test our algorithm generalization on data produced by different sources, we conducted several experiments. First, we performed leave-one-speaker-out (LOSO) using the 16 speakers of Aberystwyth English, i.e., the utterances of each speaker were tested individually with a structured model trained on the rest of the speakers. The averaged results are summarized in the first row of Table 4. Note that the results are very similar to the results shown in the second row of Table 2, where the data was randomly partitioned using 5-fold cross-validation.

Next, we evaluated our structured algorithm on a data set consisting of 607 speech intervals of the Welsh language (16 speakers from Bethesda, North Wales). The performance of the algorithm on the Welsh corpus using 4-fold cross-validation results are summarized in the second row of Table 4. Our goal was to compare the performance of the structured algorithm that was trained on one language on a different language. The last two rows in Table 4 outline the performance of the algorithm that trained on Aberystwyth English but was tested on Welsh and vice versa.

It is interesting to note that there are different performance drops when languages are mismatched. When training on the Aberystwyth data and testing on Welsh data results drop by 7.2% with a tolerance level of 10 ms. However, when tested on Aberystwyth data and training on the Welsh data we suffered only a 4.3% decrease in results, again with a 10 ms tolerance level.

Table 4: *Mismatched training and test sets. LOSO stands for leave-one-speaker-out. Thresholds and time differences in ms, values are percentages.*

Train/Test	≤5	≤10	≤15	≤20	Δt_s	Δt_e
Aber./Aber. LOSO	54.0	81.9	90.5	94.3	2.9	5.6
Welsh/Welsh	52.7	76.7	87.1	91.7	5.3	4.5
Aber./Welsh	45.8	69.5	84.4	88.8	6.6	6.5
Welsh/Aber.	54.6	80.5	90.5	94.7	3.2	5.2

5.3. Linguistic analysis

As the purpose of the algorithm is to reliably measure pre-aspiration duration for linguistic and language-related analyses, we also carried out an analysis of what language-internal and language-external factors affect pre-aspiration duration measured manually as opposed to pre-aspiration duration measured automatically with the algorithm. The outputs of four Linear Mixed Effects Models were compared. These models differed only in the dependent variable, which was (i) manually coded raw pre-aspiration duration, (ii) automatically coded raw pre-aspiration duration, (iii) manually coded normalized pre-aspiration duration, and (iv) automatically coded normalized pre-aspiration duration. The normalization method expressed




Figure 2: *Correlations between the manually coded and the predicted durations of pre-aspiration; raw (left) and normalized (right) values.*

the raw measurements as a percentage of the overall word duration.

The models shared the following independent variables: Age (continuous variable), Sex (with two levels: *female* and *male*), Vowel type (with eight levels: /a/ contrasted with /e/, /i/, /o/, /u/, /ɛ/, /ɪ/, /ɑ:/, and /ɔ:/), Place of articulation of the plosive (with three levels: /p/, /t/, /k/), and Word position (with two levels: *word-medial* and *word-final*); Speaker and Word were selected as random effects. Forward difference coding was applied to the place of articulation.

The tests show the same results regarding language-internal variables, irrespective of whether normalized or raw data are used. In all cases, pre-aspiration intervals are longest with /a/ as opposed to the other vowels ($p < 0.0001$), with the exception of /o/, which does not differ from /a/ in the durations of pre-aspiration with which it is associated. Furthermore, the duration of pre-aspiration increases as we move further back in the oral cavity ($/p/ < /t/ < /k/$; $p < 0.0001$). In addition, pre-aspiration is longer foot-finally than medially ($p < 0.0001$). The models also yield the same results concerning the external variable of sex in that no effect is found ($p = 0.06-0.39$). The only difference across the models is related to the variable of age, which comes out as significant when the dependent variable is manually coded normalized pre-aspiration (< 0.05 ; as opposed to $p = 0.06-1$). Moreover, there is a very strong positive correlation between the predicted and the manually coded values (raw measurements: $r = 0.87$; normalized measured: $r = 0.89$; $p < 0.0001$; Pearson’s product moment correlation) as shown in Figure 2.

6. Discussion

We have presented a new trainable algorithm for automatic measurement of pre-aspiration. To our knowledge, this is the first attempt to develop an automatic tool for measuring pre-aspiration. We have shown that the structured algorithm outperforms the frame-based neural network algorithm. Furthermore, we reproduced the results of a linguistic analysis based on pre-aspiration, solely using our structured algorithm’s predictions. Future work can involve extending the use of the algorithm to fricative context, automatically detecting whether there is pre-aspiration in a given interval, and applying the presented methods to lower quality data.

7. Acknowledgements

We wish to thank Jon Morris for letting us use his Welsh data.

8. References

- [1] P. Helgason, "Preaspiration in the Nordic Languages. Synchronic and Diachronic Aspects," Ph.D. dissertation, Department of Linguistics, Stockholm University, 2002.
- [2] D. Silverman, "On the rarity of pre-aspirated stops," *Journal of Linguistics*, vol. 39, pp. 575–598, 2003.
- [3] M. Clayards and T. Knowles, "Prominence enhances voicelessness and not place distinction in English voiceless sibilants," in *Proceedings of the 18th International Congress of Phonetic Sciences (ICPhS)*, Glasgow, 2015.
- [4] C. di Canio, "The phonetics of fortis and lenis consonants in Itunyoso Trique," *International Journal of American Linguistics*, vol. 78, no. 2, pp. 239–272, 2012.
- [5] G. Docherty and P. Foulkes, "Derby and Newcastle: instrumental phonetics and variationist studies," in *Urban Voices: Accent Studies in the British Isles*, P. Foulkes and G. Docherty, Eds. Routledge, 1999, pp. 47–71.
- [6] O. Gordeeva and J. Scobbie, "Preaspiration as a correlate of word-final voice in Scottish English fricatives," in *Turbulent Sounds: an Interdisciplinary Guide*, 2010, pp. 167–207.
- [7] M. Hejrná, "Preaspiration in Welsh English: A Case Study of Aberystwyth," Ph.D. dissertation, Department of Linguistics, University of Manchester, 2015.
- [8] M. Hejrná and J. Scanlon, "New laryngeal allophony in Manchester English," in *Proceedings of the 18th International Congress of Phonetic Sciences (ICPhS)*, Glasgow, 2015.
- [9] J. J. Jones and C. Llamas, "Fricated pre-aspirated /t/ in Middlesbrough English: an acoustic study," in *Proceedings of the 15th International Congress of Phonetic Sciences (ICPhS)*, Barcelona, 2003, pp. 655–658.
- [10] T. Kettig, "The BAD-LAD Split: a Phonetic Investigation," Ph.D. dissertation, Department of Theoretical and Applied Linguistics, University of Cambridge, 2015.
- [11] M. Roos, "Preaspiration in Western Yugur monosyllables," *Turgologica*, vol. 32, pp. 28–41, 1998.
- [12] M. Stevens, "How widespread is preaspiration in Italy? a preliminary acoustic phonetic overview," *Lund University Centre for Languages and Literature Phonetics Working Papers*, vol. 54, pp. 97–102, 2010.
- [13] M. Stevens and J. Hajek, "How pervasive is preaspiration? investigating sonorant devoicing in Sieneese Italian," *Proceedings of the 10th Australian International Conference on Speech Science and Technology*, pp. 334–339, 2004.
- [14] K. Watson, "The Phonetics and Phonology of Plosive Lenition in Liverpool English," Ph.D. dissertation, University of Lancaster, Edge Hill College, 2007.
- [15] A. Dwyer, "Consonantalization and obfuscation," *Turgologica*, vol. 46, pp. 423–432, 2000.
- [16] C. Ringen and W. A. van Dommelen, "Quantity and laryngeal contrasts in Norwegian," *Journal of Phonetics*, vol. 41, pp. 479–490, 2013.
- [17] C. Nance and J. Stuart-Smith, "Pre-aspiration and post-aspiration in Scottish Gaelic stop consonants," *Journal of the International Phonetic Association*, vol. 43, no. 2, pp. 129–152, 2013.
- [18] M. Hejrná, "Multiplicity of the acoustic correlates of the fortis-lenis contrast: plosives in Aberystwyth English," in *Proceeding of Interspeech*, 2016, pp. 3147–3151.
- [19] P. Helgason and C. Ringen, "Voicing and aspiration in Swedish stops," *Journal of Phonetics*, vol. 36, pp. 607–628, 2008.
- [20] P. Helgason, "The perception of medial stop contrasts in Central Standard Swedish: a pilot study," *Proceeding of FONETIK 2004, Stockholm*, pp. 92–95, 2004.
- [21] W. A. van Dommelen, "Production and perception of preaspiration in Norwegian," in *Proceeding of FONETIK 1998, Stockholm*, 1998.
- [22] M. Sonderegger and J. Keshet, "Automatic measurement of voice onset time using discriminative structured prediction," *The Journal of the Acoustical Society of America*, vol. 132, no. 6, pp. 3965–3979, 2012.
- [23] Y. Adi, J. Keshet, E. Cibelli, E. Gustafson, C. Clopper, and M. Goldrick, "Automatic measurement of vowel duration via structured prediction," *The Journal of the Acoustical Society of America*, vol. 140, no. 6, pp. 4517–4527, 2016.
- [24] Y. Adi, J. Keshet, O. Dmitrieva, and M. Goldrick, "Automatic measurement of voice onset time and prevoicing using recurrent neural networks," in *Proceeding of the 17th Annual Conference of the International Speech Communication Association (Interspeech)*, 2016.
- [25] Y. Dissen and J. Keshet, "Formant estimation and tracking using deep learning," in *Proceedings of the 17th Annual Conference of the International Speech Communication Association (Interspeech)*, 2016.
- [26] Y. Adi, J. Keshet, E. Cibelli, and M. Goldrick, "Sequence segmentation using joint RNN and structured prediction models," in *Proceeding of the 42st IEEE International Conference in Acoustic, Speech and Signal Processing (ICASSP)*, 2017.
- [27] F. Sha and L. K. Saul, "Real-time pitch determination of one or more voices by nonnegative matrix factorization," in *Proceeding of NIPS*, 2004, pp. 1233–1240.
- [28] D. Talkin, "A robust algorithm for pitch tracking (RAPT)," in *Speech coding and synthesis*, W. Kleijn and K. Paliwal, Eds. New York: Elsevier, 1995, pp. 495–518.
- [29] B. Taskar, C. Guestrin, and D. Koller, "Max-margin Markov networks," in *Proceeding of NIPS*, 2003.
- [30] I. Tsochantaridis, T. Hofmann, T. Joachims, and Y. Altun, "Support vector machine learning for interdependent and structured output spaces," in *Proceeding of the 21st International Conference on Machine Learning (ICML)*, 2004.
- [31] K. Crammer, O. Dekel, J. Keshet, S. Shalev-Shwartz, and Y. Singer, "Online passive-aggressive algorithms," *Journal of Machine Learning Research*, vol. 7, pp. 551–585, 2006.