

# Automatic measurement of vowel duration via structured prediction

Yossi Adi and Joseph Keshet

*Department of Computer Science, Bar-Ilan University, Ramat-Gan, 52900, Israel*

Emily Cibelli and Erin Gustafson

*Department of Linguistics, Northwestern University, Evanston, Illinois 60208, USA*

Cynthia Clopper

*Department of Linguistics, Ohio State University, Columbus, Ohio 43210, USA*

Matthew Goldrick<sup>a)</sup>

*Department of Linguistics, Northwestern University, Evanston, Illinois 60208, USA*

(Received 6 March 2016; revised 4 October 2016; accepted 6 December 2016; published online 27 December 2016)

A key barrier to making phonetic studies scalable and replicable is the need to rely on subjective, manual annotation. To help meet this challenge, a machine learning algorithm was developed for automatic measurement of a widely used phonetic measure: vowel duration. Manually-annotated data were used to train a model that takes as input an arbitrary length segment of the acoustic signal containing a single vowel that is preceded and followed by consonants and outputs the duration of the vowel. The model is based on the structured prediction framework. The input signal and a hypothesized set of a vowel's onset and offset are mapped to an abstract vector space by a set of acoustic feature functions. The learning algorithm is trained in this space to minimize the difference in expectations between predicted and manually-measured vowel durations. The trained model can then automatically estimate vowel durations without phonetic or orthographic transcription. Results comparing the model to three sets of manually annotated data suggest it outperformed the current gold standard for duration measurement, a hidden Markov model-based forced aligner (which requires orthographic or phonetic transcription as an input). © 2016 Acoustical Society of America. [<http://dx.doi.org/10.1121/1.4972527>]

[MAHJ]

Pages: 4517–4527

## I. INTRODUCTION

Understanding the factors that modulate vowel duration has been a long-standing focus of laboratory research in acoustic phonetics (Peterson and Lehiste, 1960). The vast majority of such research—from the mid-20th century to the start of the 21st century—has relied on manual annotation to determine vowel duration. There are two key issues with such an approach. Given the labor intensive nature of such annotation, there are substantial limitations on the amount of data that can be practically analyzed; this limits the statistical power and types of issues that phonetic studies can address. Furthermore, given the fundamental reliance on subjective annotator judgments, analyses cannot be directly replicated by other researchers.

An alternative to this approach is to pursue algorithms for automatic alignment of vowel boundaries. The current standard approach for vowels is to utilize forced alignment algorithms (Yuan *et al.*, 2013). However, this approach suffers from two important limitations: it requires an orthographic transcription, and frequently requires substantial preprocessing of the data to insure adequate performance (Evanini, 2009).

In this paper we propose a method for automatic measurement of the duration of single vowels using structured prediction techniques which have provided excellent results in analyzing other phonetic measures (Keshet *et al.*, 2007; Sonderegger and Keshet, 2012). The algorithm was trained at the segment level on manually annotated data to extract the vowel start and end times; this provides a straightforward way to compute the duration of the vowel. Following the structure of the vast majority of laboratory studies of vowel duration, we assume the input signal contains a single vowel, preceded and followed by a consonant (CVC)—no additional detailed information about the phonetic transcription is required to process the speech. This tool will allow laboratory studies to rapidly and reliably examine how the duration of vowels in monosyllabic words varies across participants and elicitation contexts.

We evaluated our method on data from three phonetic studies: one focusing on vowel duration (Heller and Goldrick, 2014; HG), the second using vowel segmentation to automatically determine points for formant analysis (Clopper and Tamati, 2014; CT), and the third a standard set of vowel production norms (Hillenbrand *et al.*, 1995; HGCW). We compared results with our model to the state-of-the-art in vowel duration measurement, HMM (hidden Markov model)-based forced alignment (Rosenfelder *et al.*, 2014). We also assessed whether inferential statistical

---

<sup>a)</sup>Electronic mail: matt-goldrick@northwestern.edu

models fit to data from our model and HMM-based forced alignment replicated the patterns obtained from manual data. The results suggest that our algorithm is superior to the current gold standard at matching the manual measurements of vowel duration, both in terms of deviation and in replicating inferential statistical results.

The paper is organized as follows. In Sec. II we state the problem definition formally. We then present the learning framework (Sec. III), algorithm (Sec. IV), and the acoustic features and feature functions (Sec. V). In Sec. VI we describe the datasets we use to train the models and evaluate the performance of the algorithm. In Sec. VII we detail the particular methods used to implement our algorithm here, along with the standard approach to vowel duration measurement. Experimental results are detailed in two sections: the first focusing on measurement deviation (Sec. VIII) and the second on the reproduction of results from inferential statistical models (Sec. IX). We conclude the paper with possible applications and extensions in Sec. X.

## II. PROBLEM SETTING

In the context of a typical laboratory study of speech, the goal of automatic vowel duration measurement is to accurately predict the time difference between the vowel onset and offset, given a segment of the acoustic signal in which the vowel is preceded by a consonant and followed by a consonant. The acoustic sample can be of any length, but should include only one vowel. We assume there is a small portion of silence before and after the uttered speech, but do not require the beginning of the speech signal or the vowel onset to be synchronized with the onset or offset of the acoustic sample.

We turn to describing the problem formally. Throughout the paper we write scalars using lowercase Latin letters, e.g.,  $x$ , and vectors using boldface letters, e.g.,  $\mathbf{x}$ . A sequence of elements is denoted with a bar  $\bar{x}$  and its length is written as  $|\bar{x}|$ . Similarly a sequence of vectors is denoted as  $\bar{\mathbf{x}}$  and its length by  $|\bar{\mathbf{x}}|$ .

The acoustic sample is represented by a sequence of acoustic feature vectors denoted by  $\bar{\mathbf{x}} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T)$ , where each  $\mathbf{x}_t$  ( $1 \leq t \leq T$ ) is a  $d$ -dimensional vector that represents the acoustic content of the  $t$ th frame. The domain of

the feature vectors is denoted as  $\mathcal{X} \subset \mathbb{R}^d$ . The input acoustic sample can be of an arbitrary length, thus  $T$  is not fixed. We denote by  $\mathcal{X}^*$  the set of all finite length segments of acoustic signal over  $\mathcal{X}$ . In addition, we denote by  $t_b \in \mathcal{T}$  and  $t_e \in \mathcal{T}$  the vowel onset and offset times in frame units, respectively, where  $\mathcal{T} = \{1, \dots, T\}$ , and the total duration of the input acoustic sample is  $T$  frames. For brevity we denote this pair by  $\mathbf{t} = (t_b, t_e)$ , and refer to it as an *onset–offset pair*. Practically there are constraints on  $t_b$  and  $t_e$  and they cannot take any value in  $\mathcal{T}$ , e.g., the vowel onset  $t_b$  cannot be  $T$  or  $T - 1$ . Our notation is depicted in Fig. 1.

Our goal is to find a function  $f$  from the domain of segments of acoustic signal,  $\mathcal{X}^*$ , to the domain of all onset–offset pairs,  $\mathcal{T}^2$ . Given a segment of the acoustic signal  $\bar{\mathbf{x}}$ , let  $\hat{\mathbf{t}} = f(\bar{\mathbf{x}})$  be the predicted onset–offset pair, where  $\hat{\mathbf{t}} = (\hat{t}_b, \hat{t}_e)$ . The quality of the prediction is assessed using a *loss function*, denote by  $\gamma(\mathbf{t}, \hat{\mathbf{t}})$ , that measures the magnitude of the penalty when predicting the pair  $\hat{\mathbf{t}}$  rather than the target pair  $\mathbf{t}$ . Formally,  $\gamma: \mathcal{T}^2 \times \mathcal{T}^2 \rightarrow \mathbb{R}_+$  is a function that receives as input two ordered pairs, and returns a positive scalar. We assume that if both the predicted and the target pairs are the same, then the loss is zero,  $\gamma(\mathbf{t}, \mathbf{t}) = 0$ .

## III. LEARNING FRAMEWORK

We assume that an input acoustic sample and a target onset–offset pair are drawn from a fixed but unknown distribution  $\rho$  over the domain of the segments of acoustic signal and the vowel onset–offset pairs,  $\mathcal{X}^* \times \mathcal{T}^2$ . We define the *risk* of  $f$  as the expected loss when using  $f$  to predict the onset–offset pair of the acoustic sample  $\bar{\mathbf{x}}$ , that is,

$$R(f) = \mathbb{E}_{(\bar{\mathbf{x}}, \mathbf{t}) \sim \rho} [\gamma(\mathbf{t}, f(\bar{\mathbf{x}}))], \quad (1)$$

where the expectation is taken with respect to the input acoustic sample  $\bar{\mathbf{x}}$  and the annotated onset–offset pair  $\mathbf{t}$  drawn from  $\rho$ . Our goal is to find  $f$  that minimizes the risk. Unfortunately, this cannot be done directly since  $\rho$  is unknown. Instead, we use a training set of examples that served as a restricted window through which we can estimate the quality of the prediction function according to the distribution of unseen examples in the real world. The examples

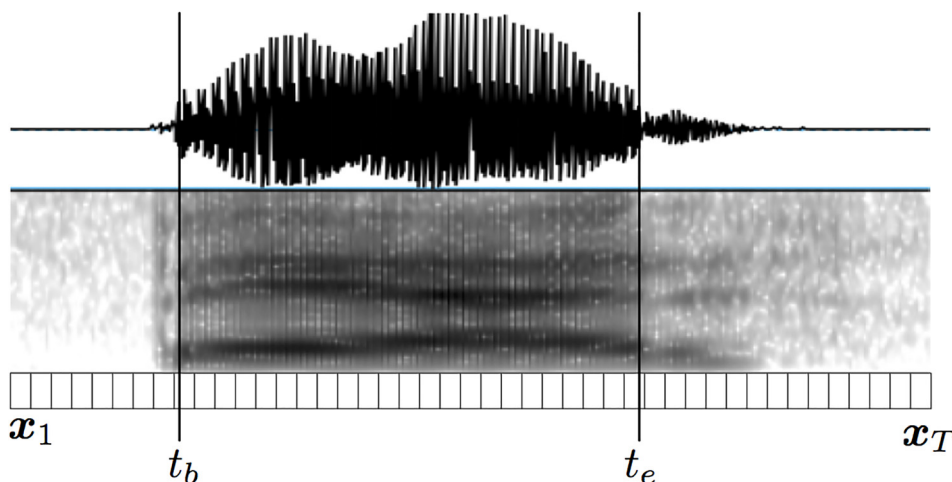


FIG. 1. (Color online) An example of our notation. The top panel presents the signal in the time domain, and the middle panel presents the spectrogram of the signal. The vertical solid lines present the annotated vowel's onset  $t_b$  and offset  $t_e$ . The speech signal is represented by a sequence of acoustic feature vectors as depicted in the lower panel. For the  $t$ th frame, the acoustic feature vector is denoted by  $\mathbf{x}_t$ .

are assumed to be identically and independently distributed according to the distribution  $\rho$ .

Each example in the training set is composed of a segment of the acoustic signal and a manually annotated vowel onset and offset pair. The manual annotations are not exact, and naturally depend both on human errors as well as objective difficulties in placing the vowel boundaries, e.g., between the vowel and a sonorant. Hence in this work we focus on a loss function that inherently takes into account the discrepancy in the annotations. That is,

$$\gamma(\mathbf{t}, \hat{\mathbf{t}}) = [|\hat{t}_b - t_b| - \tau_b]_+ + [|\hat{t}_e - t_e| - \tau_e]_+, \quad (2)$$

where  $[\pi]_+ = \max\{0, \pi\}$ , and  $\tau_b, \tau_e$  are pre-defined parameters. The above function measures the absolute differences between the predicted and target vowel onsets and offsets. It allows a mistake of  $\tau_b$  and  $\tau_e$  frames at the onset and offset of the vowel, respectively, and only penalizes predictions that are greater than  $\tau_b$  or  $\tau_e$  frames.

Our learning model belongs to the structured prediction framework. In this framework it is assumed that the output prediction is complex and has some internal structure. In our case, the vowel onset and vowel offset times are related and dependent, e.g., the vowel has a typical duration that depends on the vowel onset and offset.

In the structured prediction model, the function  $f$  is based on a fixed mapping  $\phi: \mathcal{X}^* \times \mathcal{T}^2 \rightarrow \mathbb{R}^n$  from the set of segments of acoustic signal and target onset–offset pairs to a real vector of length  $n$ ; we call the elements of this mapping *feature functions* or *feature maps*. Intuitively, the feature functions represent our knowledge regarding good locations of the onset or offset of the vowel within the acoustic signal. For example, consider Fig. 1. It can be seen that the spectral change is high at the areas of  $t_b$  and  $t_e$ . Based on this, we can adopt a feature function based on the distance between the spectrum a frame before and a frame after the presumed  $t_b$ . This function is going to be high if the presumed  $t_b$  is in the vicinity of the actual target vowel onset and is going to be low at a random place in the acoustic signal.

Our prediction function is a linear decoder with a vector of parameters  $\mathbf{w} \in \mathbb{R}^n$  that is defined as follows:

$$f_w(\bar{\mathbf{x}}) = \arg \max_{\hat{\mathbf{t}} \in \mathcal{T}^2} \mathbf{w}^\top \phi(\bar{\mathbf{x}}, \hat{\mathbf{t}}). \quad (3)$$

The subscript  $\mathbf{w}$  is added to the function  $f$  to stress that it depends on the weight vector  $\mathbf{w}$ .

Ideally, we would like our learning algorithm to find  $\mathbf{w}$  such that the prediction minimizes the loss on unseen data. Recall, we assume there exists some unknown probability distribution  $\rho$  over pairs  $(\bar{\mathbf{x}}, \mathbf{t})$ . We would like to set  $\mathbf{w}$  so as to minimize the expected loss, or the *risk*, for predicting  $f_w(\bar{\mathbf{x}})$ ,

$$\mathbf{w}^* = \arg \min_{\mathbf{w}} \mathbb{E}_{(\bar{\mathbf{x}}, \mathbf{t}) \sim \rho} [\gamma(\mathbf{t}, f_w(\bar{\mathbf{x}}))]. \quad (4)$$

It is hard to directly minimize this objective function since  $\rho$  is unknown and the loss  $\gamma$  is often a combinatorial non-

convex function (Keshet, 2014). In Sec. IV we describe the learning algorithm that aims at minimizing the risk, and then we describe the set of feature functions in Sec. V.

#### IV. DIRECT LOSS MINIMIZATION (DLM) ALGORITHM

Recall that our goal is to directly optimize the objective in Eq. (4). Unfortunately, if the output space is discrete we cannot use direct gradient descent since the loss  $\gamma(\mathbf{t}, f_w(\bar{\mathbf{x}}))$  is not a differentiable function of  $\mathbf{w}$  (Keshet, 2014). McAllester *et al.* (2010) showed that if the input space  $\mathcal{X}^*$  is continuous, we can compute the gradient of the expected loss, i.e., the risk, in Eq. (4) even when the output space is discrete in terms of the feature functions. Specifically the gradient can be expressed in a closed form solution as follows:

$$\nabla_{\mathbf{w}} \mathbb{E} [\gamma(\mathbf{t}, f_w(\bar{\mathbf{x}}))] = \lim_{\epsilon \rightarrow 0} \frac{\mathbb{E} [\phi(\bar{\mathbf{x}}, f_w^\epsilon(\bar{\mathbf{x}})) - \phi(\bar{\mathbf{x}}, f_w(\bar{\mathbf{x}}))]}{\epsilon}, \quad (5)$$

where the expectation on both sides is with respect to the tuple  $(\bar{\mathbf{x}}, \mathbf{t})$  drawn from  $\rho$ , and  $f_w^\epsilon$  is defined as follows:

$$f_w^\epsilon(\bar{\mathbf{x}}) = \arg \max_{\hat{\mathbf{t}} \in \mathcal{T}^2} \mathbf{w}^\top \phi(\bar{\mathbf{x}}, \hat{\mathbf{t}}) + \epsilon \gamma(\mathbf{t}, \hat{\mathbf{t}}). \quad (6)$$

Using stochastic gradient descent we get the following update rule:

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \frac{\eta_t}{\epsilon} \left( \phi(\bar{\mathbf{x}}, f_w(\bar{\mathbf{x}})) - \phi(\bar{\mathbf{x}}, f_w^\epsilon(\bar{\mathbf{x}})) \right), \quad (7)$$

where  $\eta_t$  is the learning rate. At training time, we set  $\eta_t = \eta_0 / \sqrt{t}$ , where  $\eta_0$  is a parameter and  $t$  is the iteration number, and we set  $\epsilon$  as a fixed small parameter, which is selected from a held-out development set. The actual values of the parameters are detailed in Sec. VIII.

Since the gradient in Eq. (5) is not a convex function in the model parameters  $\mathbf{w}$ , the gradient descent optimization method is not guaranteed to find the optimal parameter settings; it may converge to a local minimum. We initialize the model parameters with a weight vector of parameters that was pre-trained using the structured prediction passive–aggressive (PA) algorithm (Crammer *et al.*, 2006). The vector of parameters that was obtained from the PA training is denoted by  $\mathbf{w}_{\text{PA}}$ . A pseudo code of the training algorithm is given in Fig. 2.

#### V. FEATURES AND FEATURE FUNCTIONS

In this section we describe the acoustic features  $\mathbf{x} \in \mathcal{X}$  and the feature functions  $\phi(\bar{\mathbf{x}}, \mathbf{t})$ , which were designed specifically for the problem of vowel duration measurement. These features are primarily motivated by the desire to reflect the workflow of the annotators (Peterson and Lehiste, 1960). For example, in Peterson and Lehiste (1960) the beginning of the final voiced plosives was determined by comparing narrowband and broadband spectrograms in order to determine the point at which the energy in the higher harmonics is greatly diminished. In a similar fashion, we have

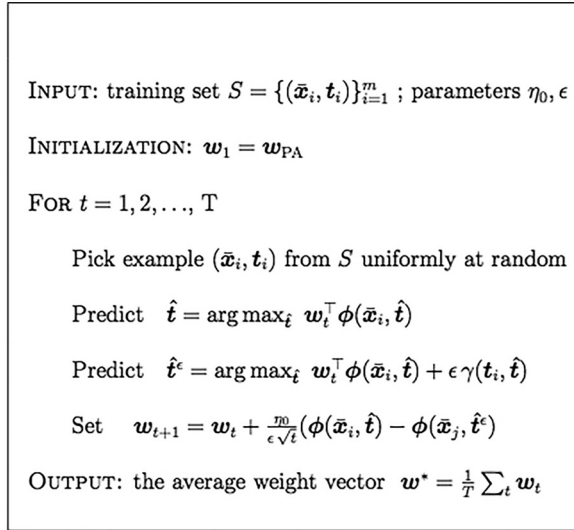


FIG. 2. DLM training procedure.

acoustic features of narrowband and broadband energies, and a set of feature functions that allow the machine learning algorithm to compare and weight them, respectively. After an empirical evaluation of performance of these features, and analyzing the source of various errors, we added features which related to the predictions of the phoneme classifiers (e.g., vowel, nasal, etc.).

### A. Acoustic features

The main function of the acoustic features representation is to preserve the crucial information-bearing elements of the speech signal and to suppress irrelevant details. We extracted  $d=16$  acoustic features every 5 ms, in a similar way to the feature set used in [Sonderegger and Keshet \(2012\)](#), but with different time spans. The first 5 features are based on the short-time Fourier transform taken with a 25 ms Hamming window. The features are short-term energy,  $E_{\text{short-term}}$ ; the log of the total spectral energy,  $E_{\text{total}}$ ; the log of the energy between 20 and 300 Hz,  $E_{\text{low}}$ ; the log of the energy above 3000 Hz,  $E_{\text{high}}$ ; and the Wiener entropy,  $H_{\text{Wiener}}$ , a measure of spectral flatness ([Sonderegger and Keshet, 2012](#)). The sixth feature,  $S_{\text{max}}$ , is the maximum of the power spectrum calculated in a region from 6 ms before to 18 ms after the frame center. The seventh feature,  $\hat{F}_0$ , is the normalized fundamental frequency estimator, extracted using the algorithm of [Sha et al. \(2004\)](#) every 5 ms, and smoothed with a Hamming window. The eighth feature is the binary output of a voicing detector based on the RAPT pitch tracker ([Talkin, 1995](#)), smoothed with a Hamming window—denoted as  $V_{\text{RAPT}}$ . The ninth feature is the number of zero crossings in a 5 ms window—denoted as  $N_{\text{ZC}}$ .

The next set of acoustic features are based on a pre-trained phoneme classifier’s predictions and scores as in [Dekel et al. \(2004\)](#). This phoneme classifier was trained on the TIMIT dataset ([Garofolo et al., 1993](#)), using the standard Mel-frequency cepstral coefficient (MFCC) features with a passive aggressive classifier ([Crammer et al., 2006](#)). The tenth acoustic feature is an estimated probability of whether

a vowel is uttered at the input frame. The feature is a smoothed version of an indicator function that states if the phoneme predicted by the classifier at the current frame is a vowel:  $\mathbf{1}[\hat{y}_t \in \text{vowels}]$ , where  $\hat{y}_t$  is the phoneme predicted by the phoneme classifier at time  $t$  and **vowels** are the set of all vowels. The 11th feature is defined to be the same as the 10th feature, but is for nasal phonemes. These features are denoted as  $G_{\text{vowel}}$  and  $G_{\text{nasal}}$ , respectively. The 12th feature is the likelihood of a vowel at the current time frame,  $L_{\text{vowel}}$ . The likelihood is computed as the Gibbs measure of the phoneme classifier’s scores for all the vowel phonemes normalized by the total score for all phonemes.

The last four features are based on the spectral changes between adjacent frames, using MFCCs to represent the spectral properties of the frames. Define by  $D_j = d(\mathbf{a}_{t-j}, \mathbf{a}_{t+j})$  the Euclidean distance between the MFCC feature vectors  $\mathbf{a}_{t-j}$  and  $\mathbf{a}_{t+j}$ , where  $\mathbf{a}_t \in \mathbb{R}^{39}$  for  $1 \leq t \leq T$ . The features are denoted by  $D_j$ , for  $j \in \{1, 2, 3, 4\}$ .

Figure 3 shows the trajectories of the features for a typical vowel in a CVC context (the word “got”).

### B. Feature functions

We turn now to describe the feature functions. Recall that the feature functions are designed to be correlated with a good positioning of the onset–offset pair,  $\mathbf{t}$ , in the acoustic signal,  $\bar{x}$ . While generally each feature function  $\phi_i(\bar{x}, \mathbf{t})$ , for  $1 \leq i \leq n$ , gets as input a sequence of acoustic features,  $\bar{x}$ , and a presumed onset–offset pair  $\mathbf{t}$ , they can practically use only a subset of the acoustic features (e.g., only a sequence over the sixth feature, as opposed to a sequence over any features). Some of the feature functions are based on the average of an acoustic feature  $x$  from frame  $t_1$  to frame  $t_2$  defined as

$$\mu(\bar{x}, t_1, t_2) = \frac{1}{t_2 - t_1 + 1} \sum_{t=t_1}^{t_2} x_t. \quad (8)$$

The functions that we identified as being correlated with a good positioning of the onset–offset pair can be divided into four groups based on the structure of the functions; each group was implemented over a set of the above acoustic features (n.b.: a given acoustic feature might be associated with multiple feature functions):

**Type 1:** This type of feature function gets as input a sequence of one of the features,  $\bar{x}$ , and a time  $t$ . The time  $t$  can represent the onset or offset of the vowel. Formally, the features of this type are of the form

$$\phi(\bar{x}, t) = x_t. \quad (9)$$

Feature functions from this type are computed for the acoustic features  $E_{\text{total}}$ ,  $E_{\text{low}}$ ,  $E_{\text{high}}$ , and  $S_{\text{max}}$  at the presumed vowel onset time  $t_b$ . It is also computed for the acoustic feature  $D_j$ ,  $j = 1, 2, 3, 4$  for both vowel onset  $t_b$  and offset times  $t_e$ . It can be seen from Fig. 3 that these acoustic features have a high value exactly at  $t_b$  or  $t_e$  or both, respectively.

**Type 2:** The second set of feature functions is a coarse estimation of the derivative around a time frame of interest.

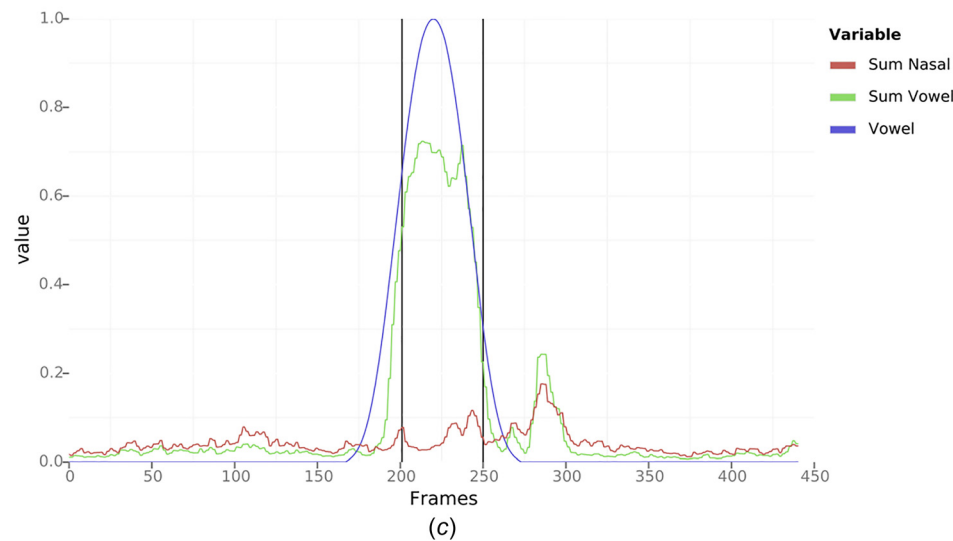
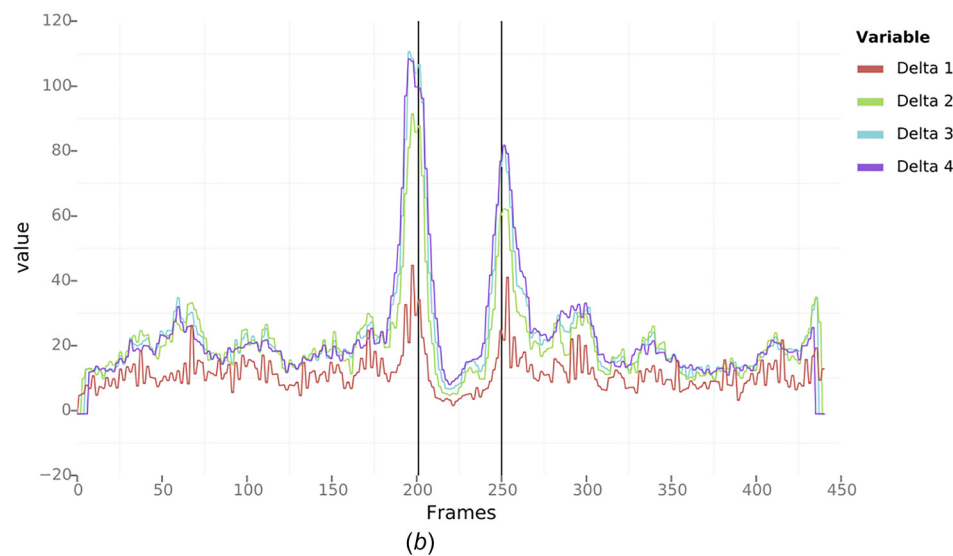
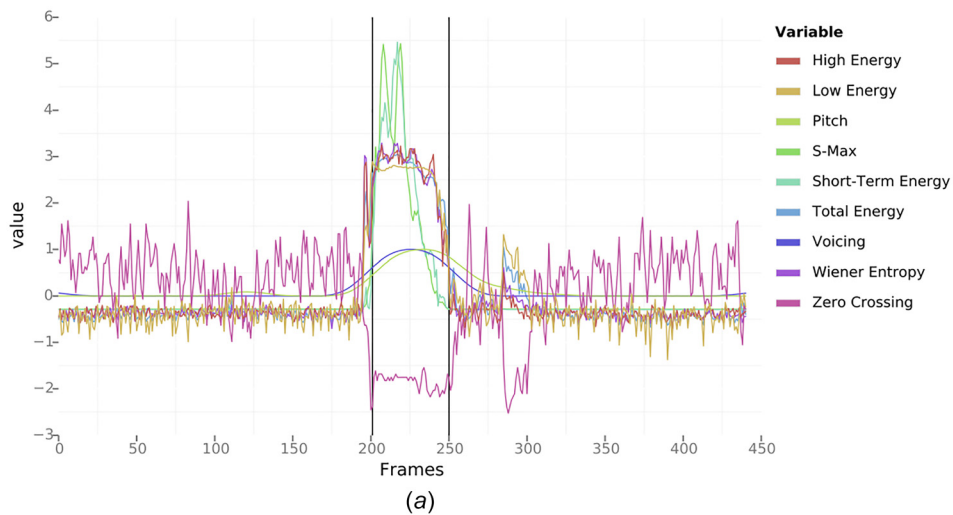


FIG. 3. (Color online) Values of acoustic features for an example acoustic sample (the word “got”). The vertical dashed lines indicate the annotated onset and offset of the vowel.

Given a sequence of acoustic feature values  $\bar{x}$  and a specific time frame  $t$ , the feature functions of this type compute the difference in the mean of  $\Delta$  frames before  $t$  and the mean of  $\Delta$  frames after  $t$ . That is,

$$\phi(\bar{x}, t, \Delta) = \mu(\bar{x}, t - \Delta, t - 1) - \mu(\bar{x}, t, t + \Delta - 1). \quad (10)$$

Table I describes all the feature functions of this type, for the values of the number of frames processed,  $\Delta$ , and the acoustic

TABLE I. The feature functions of type 2,  $\phi(\bar{x}, t_1, \Delta)$ . Values in the table are the time frame  $t_1$  used for each acoustic feature (columns) and for each window length (rows). Window lengths 8 and 10 are repeated to indicate cases where the point of interest is offset (see text for details).

$\Delta$	$E_{\text{short-term}}$	$E_{\text{low}}$	$E_{\text{high}}$	$E_{\text{total}}$	$H_{\text{Wiener}}$	$S_{\text{max}}$	$\hat{F}_0$	$V_{\text{RAPT}}$	$N_{\text{ZC}}$	$G_{\text{vowel}}$	$G_{\text{nasal}}$	$L_{\text{vowel}}$	$D_1$	$D_2$	$D_3$	$D_4$	
1						$t_b$											
2						$t_b$											
3	$t_b, t_e$					$t_b$							$t_b, t_e$	$t_b, t_e$	$t_b, t_e$	$t_b, t_e$	
4	$t_b, t_e$					$t_b$											
5	$t_b$					$t_b$											
6		$t_b$								$t_b, t_e$		$t_b, t_e$					
8		$t_b, t_e$	$t_b, t_e$	$t_b, t_e$	$t_b, t_e$		$t_b, t_e$	$t_b, t_e$	$t_b, t_e$	$t_b, t_e$	$t_e$	$t_b, t_e$					
8		$t_b - 2$	$t_b - 2$	$t_b - 2$													
10		$t_b, t_e$	$t_b, t_e$	$t_b, t_e$	$t_b, t_e$		$t_b, t_e$	$t_b, t_e$	$t_b, t_e$	$t_b, t_e$	$t_e$	$t_b, t_e$	$t_b, t_e$	$t_b, t_e$	$t_b, t_e$	$t_b, t_e$	$t_b, t_e$
10		$t_b - 4$	$t_b - 4$	$t_b - 4$													

features. For some acoustic features we wanted to take into account the possibility that they might not be coincident with vowel boundary, but occur at an adjacent point in the acoustic signal. We did that by considering the point of interest to be an offset version of the presumed onset or offset. For example,  $t_b - 2$  means the feature computes averages before and after the time frame  $t_b$  with an offset of 2 frames. It can be seen from Fig. 3 that indeed the acoustic features described in Table I have abrupt changes at the specified time frame  $t$ .

**Type 3:** The third type of feature function is an extension of the second type. Rather than considering the average of an acoustic feature before and after a single time frame, we refer to the average between the presumed onset and offset times. Feature functions of this type return two values: (i) the average between the presumed onset and offset times minus the average of  $\Delta$  frames before  $t_b$ ; and (ii) the average between the presumed onset and offset times minus the average of  $\Delta$  frames after the offset time  $t_e$ . Those two values correspond to two elements in the vector

$$\phi(\bar{x}, t_b, t_e, \Delta) = \begin{bmatrix} \mu(\mathbf{x}, t_b, t_e) - \mu(\mathbf{x}, t_b - \Delta, t_b - 1) \\ \mu(\mathbf{x}, t_b, t_e) - \mu(\mathbf{x}, t_e + 1, t_e + \Delta) \end{bmatrix}. \quad (11)$$

This type of feature function is computed for the acoustic features  $E_{\text{short-term}}$ ,  $E_{\text{low}}$ ,  $E_{\text{high}}$ ,  $E_{\text{total}}$ ,  $V_{\text{RAPT}}$ ,  $N_{\text{ZC}}$ , and  $L_{\text{vowel}}$ . Again, it can be seen from Fig. 3 that these features have high (or low) mean values between the time frames of interest.

**Type 4:** The fourth type of feature function is oblivious to the acoustic signal and returns a probability score for the presumed vowel duration,  $t_e - t_b$ . We have two feature functions of this type. The first one is based on the Normal distribution with parameters  $\hat{\mu}$ ,  $\hat{\sigma}^2$  that are estimated from the training data

$$\phi(t_b, t_e) = \mathcal{N}(t_e - t_b; \hat{\mu}, \hat{\sigma}^2). \quad (12)$$

Similarly, the second feature function of this form is based on the Gamma distribution with parameters  $\hat{k}$  and  $\hat{\theta}$ , that are estimated from the training set

$$\phi(t_b, t_e) = \Gamma(t_e - t_b; \hat{k}, \hat{\theta}). \quad (13)$$

While the Gamma distribution is an appropriate distribution to describe the vowel duration alone, we found empirically that adding the Normal distribution improves performance.

## VI. DATASETS

In order to get reliable results we used three different datasets to evaluate the performance of our system. In this section we will give a short description of each dataset.

### A. HG

This corpus (HG) is drawn from a study investigating how grammatical class constraints influence the activation of phonological neighbors. The influence of neighbors on phonetic processing was indexed by vowel durations. It contains segments of acoustic signal from 64 native English speakers (55 female) aged 18–34 with no history of speech or language deficits. Participants were first familiarized to a set of pictures along with their intended labels. They were then asked to name aloud the noun depicted by a picture in two contexts: when the picture was presented alone, and when the picture occurred at the end of a non-predictive sentence. Participants were instructed to produce the name as quickly and accurately as possible. Trials with errors or disfluencies were excluded, along with two items with high error rates. In addition, data from three subjects reported in the original paper were excluded from the present analysis, as they did not consent to public use of their data. The remaining 2395 recorded segments of acoustic signal contain one English CVC noun with vowels /i, ε, ae, α, oʊ, u/.

### B. CT

The second dataset (CT) contains segments of acoustic signal from 20 female native English speakers aged 18–22 with no history of speech or language deficits. The participants were evenly split between two American English dialects (Northern and Midland). As part of a larger study (Clopper *et al.*, 2002), participants read aloud a list of 991 CVC words. This study focused on 39 target words (777 tokens) which did or did not have a lexical contrast between either /ε/ vs /ae/ (e.g., dead-dad vs deaf-\*daff) or /α/ vs /ɔ/ (e.g., cot-caught vs dock-\*dawk). Words with a lexical

contrast are referred to as *competitor* items and those without are referred to as *no competitor* items.

### C. HGCW

The third dataset consists of data from a laboratory study conducted by HGCW. It contains segments of acoustic signal from 45 men, 48 women, and 46 10- to 12-yr-old children (27 boys and 19 girls). The participants (87%) were raised in Michigan, primarily in the southeastern and southwestern parts of the state. The audio recordings contain 12 different vowels (/i, ɪ, ε, ae, ɑ, ɔ, ʊ, u, ʌ, ɜ, e, o/) from the words: heed, hid, head, had, hod, hawed, hood, who would, hud, heard, hayed, hoed.

## VII. METHODS

Code implementing this algorithm is publicly available at <https://github.com/adiyoss/AutoVowelDuration>. Segments of acoustic signal for the HG dataset along with algorithmic and manual annotations are available in the Online Speech/Corpora Archive and Analysis Resource (<https://oscaar.ci-northwestern.edu/>; dataset “Within Category Neighborhood Density”). For the HGCW dataset, segments of acoustic signal and manual annotations are available at <http://homepages.wmich.edu/~hillenbr/voweldata.html>.

### A. DLM

The DLM algorithm described above was trained and tested on both HG and CT datasets. For the CT corpus, the training set was a set of 4049 tokens from the original corpus (Clopper *et al.*, 2002), and the test set was the same 777 tokens reported in CT. The model parameters were tuned on a dedicated development set (~10% of the training set). The parameters that yielded the lowest error rate were  $\eta = 0.1$  and  $\epsilon = -1.36$ . We used  $\tau_b = 1$  frame and  $\tau_e = 2$  frames for the loss during training. The initial weight vector was set to be the averaged weight vector from the PA algorithm (Crammer *et al.*, 2006) with  $C = 0.5$  and 100 epochs. When trained and tested on the HG corpus, we used tenfold cross validation (the reported results are the average error over all tenfolds) with the same settings and parameters as described above.

In order to better comprehend the influence of the phoneme classifier, the only language dependent feature in our system, we also trained and tested the direct loss algorithm without the acoustic features related to the phoneme classifier  $G_{\text{vowel}}$ ,  $G_{\text{nasal}}$ , or  $L_{\text{vowel}}$ , and their corresponding feature functions. In the case when the phoneme classifier was used, we trained multiclass PA as described in Dekel *et al.* (2004) on the TIMIT corpus of read speech.

### B. HMM

To validate the effectiveness of the proposed approach, we compared it to the most common approach currently used in automatic phonetic measurement of speech: forced alignment. This is an algorithm which, given a speech utterance and its phonetic content, finds the start time of each phoneme in the speech utterance. Often the orthographic

content of the speech utterance is given, and it is converted to its phonetic content using a lexicon. This procedure may not be accurate as often the surface pronunciation uttered in spontaneous speech is not the same as the canonical pronunciation that appears in the lexicon.

Conventionally, the forced alignment is implemented by forcing the decoder of an HMM phoneme recognizer to pass through the states corresponding to the given phoneme sequence. So far, automatic vowel extraction has been done using such forced alignment procedures (Reddy and Stanford, 2015). While there are several open source implementations of HMMs, all the publicly available forced aligner packages (Goldman, 2011; Gorman *et al.*, 2011; Rosenfelder *et al.*, 2014; Yuan and Liberman, 2008) are based on the HTK toolkit (Young and Young, 1994). Since this is the case, we chose to compare our results with the most recent one, namely, the *FAVE aligner* (Rosenfelder *et al.*, 2014), based on the *Penn Phonetics Lab Forced Aligner* (Yuan and Liberman, 2008). In our analyses below, we refer to this system as the HMM aligner. Note that in contrast to the DLM the forced aligner requires a phonetic or orthographic transcription as input.

Since neither CT nor HG has enough data to train an HMM forced alignment system, in this work we use a pre-trained model which was trained on different corpus [see details in Rosenfelder *et al.* (2014)]. In order to make a fair comparison between the proposed model and the HMM system, we evaluate the models on two different settings: (i) a dataset which was not used during training any of the models (HGCW), and (ii) on mismatched training and test datasets.

## VIII. RESULTS: MEASUREMENT DEVIATION

We first examine results on the HG and CT datasets when the DLM algorithm is trained and tested on the same corpus. The difference between the automatic and manual measurements of vowel duration are given in Table II, where the evaluation metric is the loss in Eq. (2) with both  $\tau_b$  and  $\tau_e$  equal to 0, and in Table III, where the error is given in terms of the percentage of predictions that do not fall within the boundaries of 20 ms from the manual onset and 50 ms from the manual offset. (We allow for greater deviations in vowel offset, where manual annotators often show a greater disagreement.) These reveal that when trained and tested on the same corpus, the DLM generally outperforms the standard HMM approach, particularly when the phoneme classifier is incorporated into the algorithm. An exception is seen

TABLE II. Results of DLM (with and without phoneme classifier, trained and tested on the same corpus) and HMM relative to manual annotation. Average deviation of onset and offset (in msec). Bold indicates minimum deviation for each dataset within each context.

	DLM		DLM (no classifier)		HMM	
	onset	offset	onset	offset	onset	offset
HG	<b>5.21</b>	<b>22.80</b>	5.84	28.98	16.67	24.72
CT	9.42	<b>16.76</b>	<b>9.24</b>	23.18	35.90	30.61

TABLE III. The percentage of predictions that do not fall within the boundaries of 20 ms at the onset and 50 ms at the offset from the manual annotation when the DLM is trained and tested on the same corpus. Bold indicates minimum deviation for each dataset within each context.

	DLM		DLM (no classifier)		HMM	
	onset	offset	onset	offset	onset	offset
<i>HG</i>	<b>6.15%</b>	13.15%	8.05%	18.44%	31.90%	<b>10.94%</b>
<i>CT</i>	<b>9.46%</b>	<b>8.31%</b>	9.59%	9.34%	41.50%	13.14%

in the HG offset data, where the percentage of predictions outside of 50 ms is smallest for the HMM.

It is possible that the accuracy of the algorithms may differ by consonantal context. To assess this, we classified the consonant preceding or following the vowel of each word as belonging to one of four categories: voiceless stops, voiced stops, fricatives/affricates, and sonorants. We then calculated the overall measurement error for each algorithm within each category, as reported in Table IV.

The context-specific HG data shows that the general advantage for the DLM with the phoneme classifier holds, but there are some exceptions. The performance of the HMM is comparable to the DLM with classifier when a vowel is followed by a voiceless stop, and appears to outperform the DLM when voiceless stops or fricatives precede the vowel. In the CT data, there is a clear DLM advantage across all contexts.

### A. Mismatched training and test datasets

The comparison above may be biased in favor of the DLM, as the algorithm’s parameters may be tuned to specific features of the annotators that worked with the HG and CT corpora. To provide a more even comparison, the DLM was trained and tested on different corpora—training on HG and testing on CT and vice versa. We also present the result of our algorithm trained on CT (the manual dataset with the

TABLE IV. Results of DLM (with and without phoneme classifier, trained and tested on the same corpus) and HMM relative to manual annotation, broken down by consonant context. Onset consonant context is shown in the first, followed by coda consonant context. Average deviation of full vowel duration [in msec]. Fricatives are all voiceless, with the exception of [dʒ] in CT coda consonants. There were no sonorants in CT codas. Bold indicates minimum deviation within each context for each dataset.

	DLM		DLM (no classifier)		HMM	
	HG	CT	HG	CT	HG	CT
<i>Onset</i>						
<i>Voiceless stops</i>	25.67	17.07	30.65	<b>16.96</b>	<b>21.95</b>	97.45
<i>Voiced stops</i>	<b>27.81</b>	<b>18.62</b>	31.23	20.27	31.66	53.90
<i>Fricatives/affricates</i>	25.03	16.74	32.40	<b>16.28</b>	<b>18.01</b>	32.86
<i>Sonorants</i>	<b>24.12</b>	<b>26.97</b>	36.32	27.95	42.35	40.52
	DLM		DLM (no classifier)		HMM	
<i>Coda</i>	HG	CT	HG	CT	HG	CT
<i>Voiceless stops</i>	<b>25.72</b>	18.06	28.76	<b>17.78</b>	25.91	43.22
<i>Voiced stops</i>	<b>16.16</b>	<b>23.08</b>	20.10	27.22	19.04	45.19
<i>Fricatives/affricates</i>	<b>15.00</b>	20.42	27.25	<b>20.41</b>	27.92	72.82
<i>Sonorants</i>	<b>32.51</b>	—	39.04	—	35.36	—

highest manual inter-annotator reliability; see below) and tested on the third dataset, HGCW.

As before, the difference between the automatic and manual measurements of vowel duration are given in Table V, and the error in terms of the percentage of predictions that do not fall within the boundaries of 20 ms from the manual onset and 50 ms from the manual offset is given in Table VI. These results again show that for the onset data, the DLM outperforms the HMM, even when the training and testing sets for the algorithm are mismatched. In vowel offsets, the HMM meets or outperforms the DLM for the HG and CT data, but the DLM advantage is maintained for the HGCW offset data.

### B. Measurement correlation

Above, we examined the degree to which each algorithm agreed with manual annotators based on the amount of deviation between measures. Another measure of agreement conventionally reported in phonetic studies is the correlation between measurements of vowel durations. This is typically done by assigning a random subset of the data to a second annotator. For the HG dataset, the second annotator’s correlation with the original annotator was  $r(627) = 0.84$ . We performed a similar analysis, treating the algorithm as a second annotator relative to the manual labels. For the same subset of the data (with the DLM trained and tested on the same corpus), the correlations between each algorithm and the original manual annotator are DLM = 0.79, DLM (no classifier) = 0.64, and HMM = 0.73. While using the models that were trained on mismatched data the correlations are DLM = 0.67, DLM (no classifier) = 0.55. Similarly, for the CT dataset, the second annotator’s correlation with the original annotator was  $r(398) = 0.95$ . For the same subset of the data, each algorithm’s correlations are: DLM = 0.93, DLM (no classifier) = 0.93, and HMM = 0.57. However, when using the models that were trained on mismatched data the correlations are DLM = 0.54, DLM (no classifier) = 0.53. Thus, when training the models on mismatched datasets, correlations of the HMM model and DLM model with human annotators are equivalent; the DLM only outperforms HMM when using the same training and testing set.

### C. Discussion

Analysis of measurement deviation suggests our model generally outperforms the HMM-forced alignment algorithm, even when training and testing sets are mismatched.

TABLE V. Results of DLM (with and without phoneme classifier) and HMM under mismatched training and test datasets. Average deviation of onset and offset [in msec]. Bold indicates minimum deviation for each dataset within each context.

	DLM		DLM (no classifier)		HMM	
	onset	offset	onset	offset	onset	offset
<i>HG (CT model)</i>	<b>14.44</b>	38.66	14.92	43.95	16.67	<b>24.72</b>
<i>CT (HG model)</i>	<b>9.84</b>	30.85	10.17	<b>29.94</b>	35.90	30.61
<i>HGCW (CT model)</i>	<b>12.91</b>	<b>9.28</b>	15.86	23.89	19.95	27.30



TABLE VI. The percentage of predictions that do not fall within the boundaries of 20 ms at the onset and 50 ms at the offset from the manual annotation (with mismatched training and test datasets for the DLM). Bold indicates minimum deviation for each dataset within each context.

	DLM		DLM (no classifier)		HMM	
	onset	offset	onset	offset	onset	offset
HG (CT model)	<b>11.55%</b>	23.99%	12.25%	28.80%	31.90%	<b>10.94%</b>
CT (HG model)	<b>8.36%</b>	14.92%	<b>8.36%</b>	<b>10.94%</b>	41.50%	13.14%
HGCW (CT model)	<b>19.18%</b>	<b>2.16%</b>	27.58%	6.54%	28.54%	11.69%

Incorporating the phoneme classifier improved performance, but even without the classifier the DLM typically outperforms the HMM. However, with respect to measurement correlation, there was no clear advantage; DLM and HMM exhibited equivalent performance when trained on mismatched training and test sets.

## IX. RESULTS: REPRODUCING REGRESSION MODEL FINDINGS

Empirical studies of speech and language processing use acoustic properties such as vowel duration as behavioral measures of the effects of various types of variables that influence speech and language. Canonical examples of variables of study include the phonetic context of vowels (e.g., preceding voiced vs voiceless stops), properties of speakers who produce those vowels (e.g., native vs non-native speakers), and the linguistic (lexical, syntactic, etc.) context in which the vowels appear (e.g., predictable vs unpredictable). The effects of these variables on acoustic properties are typically examined using inferential statistical models.

The second evaluation metric of our algorithms therefore examined the similarity of inferential statistical model fits based on measurements generated by algorithms vs manual measurements. Out of the three datasets, only HG constructed models based on duration data; our analysis therefore focused on this dataset.<sup>1</sup>

HG examined whether processes involved in sentence planning influence the processing of sound. Speakers named pictures in a context that strongly emphasized sentence planning (following a sentence fragment) as well as a context that did not require substantial planning (producing the picture name in isolation). The order of these contexts was counterbalanced across speakers. To index effects of sound structure processing, this study also manipulated the number of words phonologically similar to the target that share its grammatical category (category-specific lexical density).

To analyze the effect of this variable, HG used linear mixed effects regression models (Baayen *et al.*, 2008), an approach that has become dominant in the analysis of speech and psycholinguistic data. These regression models predict dependent measures based on a linear combination of predictors, including both fixed and randomly distributed predictors capturing variation in effects by both participants and items. This allows researchers to examine the effects of interest while controlling for other properties of the words and speakers.

Analysis of the full data set showed that lexical density and vowel category had no effect on vowel durations (Heller and Goldrick, 2014, 2015); however, these effects were not stable when the subset of data used here was examined. Therefore, the models used here were simplified from those reported in the original paper, including two contrast-coded fixed effect factors: production context (isolation vs sentence) and block (first vs second). To control for contributions from the random sample of participants and items used in this experiment (compared to all English speakers and all words of English), we included two sets of random effects. Random intercepts for both speakers and words were included, along with uncorrelated slopes for context by both speaker and word.

## A. Significance of fixed effects in models of the HG dataset

We first examine the primary interest of many phonetic studies—the binary distinction between significant vs insignificant effects of fixed-effect predictors (e.g., the effect of experimental condition). To control for skew, vowel durations were log-transformed prior to analysis. Given that outliers can have an outsized influence on parameter estimates, all regressions models were refit after excluding observations with standardized residuals exceeding 2.5 (Baayen, 2008). The significance of fixed-effects predictors was assessed by using the likelihood ratio test to compare models with and without the predictor of interest (Barr *et al.*, 2013). Table VII compares the estimates for the two fixed effects parameters for the manual model and each algorithmic model. Although there was some variation in the model estimates of the fixed effects for each algorithm compared to the manual annotations, each algorithm recovered the overall pattern of a significant effect of context but not block.

### 1. Comparison of predictions of models of the HG dataset

An alternative, more global, assessment of model similarity is to compare model predictions. This was assessed by leave-one-out validation. For each observation, we excluded it from the dataset and re-fit the regression to the remaining observations. After residual-based outlier trimming (and re-fitting), we examined the predictions of this re-fitted model for the excluded observation. Figure 4 shows distributions of the deviation of each algorithm’s predicted fit to the predicted model fit to the manual data.

The DLM algorithm outperformed the other algorithms. The mean squared error relative to the predictions of the

TABLE VII. Estimates and *t*-values for fixed effects in regression model, HG dataset (standard error of estimate in parentheses). Significant effects ( $p < 0.05$ , as assessed by likelihood ratio tests) are bolded.

	Context	<i>t</i>	Block	<i>t</i>
Manual	<b>-0.057 (0.017)</b>	<b>-3.29</b>	0.012 (0.016)	0.72
DLM	<b>-0.072 (0.019)</b>	<b>-3.82</b>	0.016 (0.018)	0.91
DLM (no classifier)	<b>-0.059 (0.018)</b>	<b>-3.34</b>	0.015 (0.017)	0.90
HMM	<b>-0.058 (0.018)</b>	<b>-3.13</b>	0.008 (0.014)	0.55

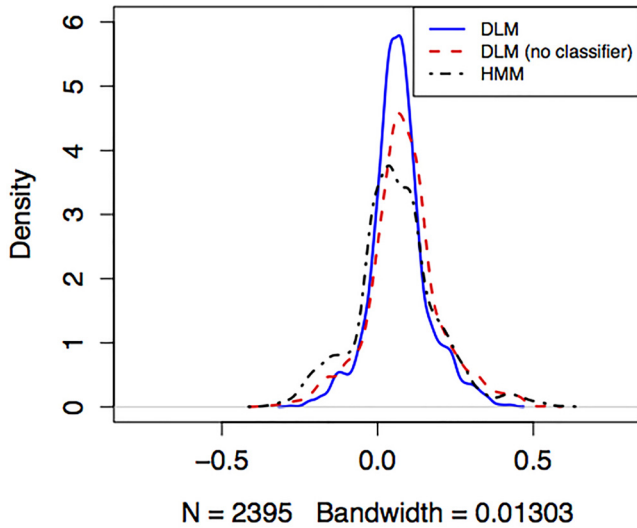


FIG. 4. (Color online) Comparisons of leave-one-out model predictions of vowel durations for the HG dataset. Deviance of each algorithmic method as compared to the manual predictions (manual-algorithmic) are plotted as individual lines.

model fit to the manually annotated data was lowest for the DLM algorithm (0.479 msec<sup>2</sup>), higher for HMM (0.835 msec<sup>2</sup>), and larger still for DLM (no classifier) (0.935 msec<sup>2</sup>). Bootstrap confidence intervals (CIs) of the differences across algorithms showed that DLM outperformed the other two algorithms (95% CI of differences from HMM: [0.00039, 0.00043] msec<sup>2</sup>; DLM (no classifier): [0.00037, 0.00055] msec<sup>2</sup>). The HMM algorithm did not significantly differ from DLM (no classifier) (95% CI [-0.000200, 0.000002] msec<sup>2</sup>).

Parallel to the analysis of measurement deviation, these analyses suggest our model outperforms the HMM-forced aligner. Incorporating the phoneme classifier improved performance, but even without the classifier the performance of the DLM matches that of the HMM.

## X. GENERAL DISCUSSION

We presented an algorithm for automatically estimating the durations of vowels based on the structured prediction framework, relying on a set of acoustic features and feature functions finely tuned to reflect the properties of vowel acoustics relevant to vowel segmentation. The DLM algorithm, when including a phoneme classifier, clearly succeeds at matching manual measurements of vowel duration. With respect to both measurement deviation and reproduction of regression model results, the DLM algorithm outperforms or matches a commonly-used forced alignment system while not requiring a phonetic transcription. This approach can allow laboratory experiments to address much larger samples of data in a way that is replicable and reliable.

Having achieved some success with monosyllabic, laboratory stimuli, future development of this algorithm should extend this approach to more naturalistic production. In current work, we are extending this approach to vowel durations in multisyllabic words. We believe that moving outside the laboratory will be facilitated by the structure of our

approach; in contrast with existing systems, our algorithm has the additional benefit of not requiring a transcript of the desired speech prior to analysis. Extending the capability of the system to analyze untranscribed speech will support the analysis of more naturalistic, connected speech, including speech styles or dialects that may be difficult to robustly sample in a laboratory context (Labov, 1972; Rischel, 1992).

With respect to the algorithms utilized in future work, we aim to better determine the acoustic features that allow the most precise estimation of vowel duration. While our current work relies on features determined by expert annotators (Peterson and Lehiste, 1960), in ongoing work we aim to have the algorithm discover these (as well as perhaps novel features not used by experts). We are currently pursuing this by applying new advances in deep learning to automatic estimation of vowel duration. In preliminary work, we used this approach to train a binary classifier to detect for each frame whether it contains a vowel or not (Adi *et al.*, 2015). However, since this approach does not take into account the relationships between adjacent frames, it requires an additional algorithm to yield durations. In future work, we plan to focus on combining new advances in sequence deep learning (Elman, 1990; Graves *et al.*, 2013; Graves and Jaitly, 2014) with our current structured prediction scheme.

## ACKNOWLEDGMENTS

Research supported by NIH Grant No. 1R21HD077140 and NSF Grant No. BCS1056409.

<sup>1</sup>See supplementary material at <http://dx.doi.org/10.1121/1.4972527> for regression analysis of the CT dataset, where the vowel duration measurements of our algorithm as well as the HMM are provided as inputs to a formant estimator.

- Adi, Y., Keshet, J., and Goldrick, M. (2015). "Vowel duration measurement using deep neural networks," in *Proceedings of the IEEE International Workshop on Machine Learning for Signal Processing*, pp. 1–6.
- Baayen, R. H. (2008). *Analyzing Linguistic Data: A Practical Introduction to Statistics using R* (Cambridge University Press, Cambridge).
- Baayen, R. H., Davidson, D. J., and Bates, D. M. (2008). "Mixed-effects modeling with crossed random effects for subjects and items," *J. Memory Lang.* **59**(4), 390–412.
- Barr, D. J., Levy, R., Scheepers, C., and Tily, H. J. (2013). "Random effects structure for confirmatory hypothesis testing: Keep it maximal," *J. Memory Lang.* **68**(3), 255–278.
- Clopper, C. G., Carter, A. K., Dillon, C. M., Hernandez, L. R., Pisoni, D. B., Clarke, C. M., Harnsberger, J. D., and Herman, R. (2002). "The Indiana speech project: An overview of the development of a multi-talker multi-dialect speech corpus," Indiana University Speech Research Laboratory Research on Speech Perception Progress Report, Vol. 25, pp. 367–380.
- Clopper, C. G., and Tamati, T. N. (2014). "Effects of local lexical competition and regional dialect on vowel production," *J. Acoust. Soc. Am.* **136**(1), 1–4.
- Crammer, K., Dekel, O., Keshet, J., Shalev-Shwartz, S., and Singer, Y. (2006). "Online passive aggressive algorithms," *J. Machine Learn. Res.* **7**, 551–585.
- Dekel, O., Keshet, J., and Singer, Y. (2004). "An online algorithm for hierarchical phoneme classification," in *Workshop on Multimodal Interaction and Related Machine Learning Algorithms; Lecture Notes in Computer Science*, Vol. 3361/2005 (Springer-Verlag, Berlin), pp. 146–159.
- Elman, J. L. (1990). "Finding structure in time," *Cognitive Sci.* **14**(2), 179–211.

- Evanini, K. (2009). "The permeability of dialect boundaries: A case study of the region surrounding Erie," Ph.D. thesis, University of Pennsylvania.
- Garofolo, J. S., Lamel, L. F., Fisher, W. M., Fiscus, J. G., Pallett, D. S., Dahlgren, N. L., and Zue, V. (1993). "TIMIT acoustic-phonetic continuous speech corpus," in *Linguistic Data Consortium*, Philadelphia, PA, 33 pp.
- Goldman, J.-P. (2011). "Easyalign: An automatic phonetic alignment tool under Praat," in *Proceedings of Interspeech*.
- Gorman, K., Howell, J., and Wagner, M. (2011). "Prosodylab-aligner: A tool for forced alignment of laboratory speech," *Canadian Acoust.* **39**(3), 192–193.
- Graves, A., and Jaitly, N. (2014). "Towards end-to-end speech recognition with recurrent neural networks," in *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pp. 1764–1772.
- Graves, A., Mohamed, A.-R., and Hinton, G. (2013). "Speech recognition with deep recurrent neural networks," in *Proceedings of 2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6645–6649.
- Heller, J. R., and Goldrick, M. (2014). "Grammatical constraints on phonological encoding in speech production," *Psychonomic Bull. Rev.* **21**(6), 1576–1582.
- Heller, J. R., and Goldrick, M. (2015). "Erratum to: 'Grammatical constraints on phonological encoding in speech production,'" *Psychonomic Bull. Rev.* **22**(5), 1475.
- Hillenbrand, J., Getty, L. A., Clark, M. J., and Wheeler, K. (1995). "Acoustic characteristics of American English vowels," *J. Acoust. Soc. Am.* **97**(5), 3099–3111.
- Keshet, J. (2014). "Optimizing the measure of performance in structured prediction," in *Advanced Structured Prediction*, edited by S. Nowozin, P. V. Gehler, J. Jancsary, and C. H. Lampert (MIT Press, Cambridge, MA).
- Keshet, J., Shalev-Shwartz, S., Singer, Y., and Chazan, D. (2007). "A large margin algorithm for speech-to-phoneme and music-to-score alignment," *IEEE Trans. Audio, Speech Lang. Process* **15**, 2373–2382.
- Labov, W. (1972). *Sociolinguistic Patterns* (University of Pennsylvania Press, Philadelphia, PA).
- McAllester, D., Hazan, T., and Keshet, J. (2010). "Direct loss minimization for structured prediction," in *Proceedings of the 23rd International Conference on Neural Information Processing Systems*, Vancouver, British Columbia, Canada (December 6–9, 2010).
- Peterson, G. E., and Lehiste, I. (1960). "Duration of syllable nuclei in English," *J. Acoust. Soc. Am.* **32**(6), 693–703.
- Reddy, S., and Stanford, J. N. (2015). "Toward completely automated vowel extraction: Introducing DARLA," *Linguistics Vanguard* **1**, 15–28.
- Rischel, J. (1992). "Formal linguistics and real speech," *Speech Commun.* **11**(4), 379–392.
- Rosenfelder, I., Fruehwald, J., Evanini, K., Seyfarth, S., Gorman, K., Prichard, H., and Yuan, J. (2014). "Fave (forced alignment and vowel extraction)," Program suite v1.2.2 10.5281/zenodo.22281.
- Sha, F., Burgoyne, J. A., and Saul, L. K. (2004). "Multiband statistical learning for f0 estimation in speech," in *Proceeding of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Vol. 5, pp. 661–664.
- Sonderegger, M., and Keshet, J. (2012). "Automatic discriminative measurement of voice onset time," *J. Acoust. Soc. Am.* **132**(6), 3965–3979.
- Talkin, D. (1995). "A robust algorithm for pitch tracking," in *Speech Coding and Synthesis*, edited by W. B. Kleijn and K. K. Paliwal (Elsevier, New York), pp. 495–518.
- Young, S., and Young, S. (1994). "The HTK hidden Markov model toolkit: Design and philosophy," *Entropic Cambridge Res. Lab., Ltd.* **2**, 2–44.
- Yuan, J., and Liberman, M. (2008). "Speaker identification on the SCOTUS corpus," *J. Acoust. Soc. Am.* **123**(5), 3878.
- Yuan, J., Ryant, N., Liberman, M., Stolcke, A., Mitra, V., and Wang, W. (2013). "Automatic phonetic segmentation using boundary models," in *Proceedings of Interspeech*, pp. 2306–2310.