Extracting dynamic parameters from speech movement data

858

Caroline L. Smith[a)]
Haskins Laboratories, 270 Crown Street, New Haven, Connecticut 06511 and Department of Linguistics, Yale University, New Haven, Connecticut 06520

Catherine P. Browman and Richard S. McGowan
Haskins Laboratories, 270 Crown Street, New Haven, Connecticut 06511

Bruce Kay
Department of Psychology, Brown University, Providence, Rhode Island 02912

(Received 15 November 1991; accepted for publication 30 November 1992)

A quantitative characterization of articulatory movements, using the parameter values of a linear second-order dynamical system, was developed in order to compare classes of movements, in particular, classes defined by linguistic factors such as syllable position, stress, and vowel quality. Movements of the lower lip in utterances such as [ˈbibəbib] and [babəˈbab] were partitioned into sections ("windows") in two ways: at successive displacement peaks and valleys, and at the right edge of plateau regions around such extreme values. The linguistic factors affected natural frequency in similar ways regardless of whether damping ratio was permitted to vary or held fixed at one of several different values. Damping ratio was generally unaffected by the linguistic factors. For the most part, the type of partition or window did not affect the patterns of the results, with the exception of the closing gesture out of the reduced syllable.

PACS numbers: 43.70.Aj, 43.70.Bk, 43.72.Ar

INTRODUCTION

Analysis of articulatory movement data requires a quantitative description of the spatial and temporal properties of the movement. A useful description has fewer degrees of freedom than the data, and captures movement characteristics due to membership in a class while also accurately representing idiosyncratic properties. Such a description should also facilitate comparison among classes of related movements. Different researchers have used different techniques to describe movement data; we will be reporting on a technique within the framework of dynamical system theory (e.g., Sonoda and Kiritani, 1976; Fowler *et al.*, 1980; Ostry *et al.*, 1983; Ostry and Munhall, 1985; Browman and Goldstein, 1985; Kelso *et al.*, 1986; Saltzman and Munhall, 1989; Perrier *et al.*, 1991).

The equations of motion for a dynamical system represent changes in spatial variables over time by stating a relationship among variables of motion that remains constant over time. An example of a simple dynamical system is the mass–spring system, whose response to forces acting upon it can be expressed in terms of a linear second-order differential equation. This type of dynamical system has been shown to produce time series for articulator displacements with that connection between displacement and peak velocity that is characteristic of (reiterant) speech movements over changes in stress and speaking rate (Tuller *et al.*, 1982; Ostry and Munhall, 1985; Kelso *et al.*, 1985; Vatikiotis-Bateson, 1988). In particular, changes in stress have been treated as changes in articulator stiffness (Munhall *et al.*, 1985; Kelso *et al.*, 1985; Browman and Goldstein, 1985), since change in the displacement/peak velocity relationship can be modeled in a second-order system as a change in articulator stiffness (Cooke, 1980).

Modeling articulatory data as a dynamical system relates the data to a well-defined system with a constrained description. Thus a relatively small number of parameters is needed to specify a particular time series within a given system. An overall comparison of two classes of movements can be made by comparing the two sets of parameter values that represent the classes rather than comparing two sets of time series directly. Note that we are not intending here to compare a damped mass-spring model with other dynamic models of speech articulation, such as that of Perrier *et al.* (1989). Rather, the work reported here investigates how the parameter values for a second-order dynamical system, particularly stiffness and damping, reflect changes in linguistic factors such as stress, syllable position, and vowel quality. We further investigated the consequences of holding damping ratio constant for the patterns found in stiffness.

Finally, we investigated the consequences for the parameter values of using different methods of sectioning the movement curves into "windows." That is, a particular set of parameter values describes a particular dynamical system that fits the articulatory data for a period of time. But the system state controlling the articulatory movements does not remain the same indefinitely: continuous movement is associated with a set of phonetically discrete units, or "gestures" (see Browman and Goldstein, 1989). The values of the system parameters change between different phonetic units. Thus the movements themselves must be divided up into the sections, i.e., windows, that correspond to different

[a)] Now at Dept. of Linguistics, UCLA, Los Angeles, CA 90024.

control regimes for discrete phonetic units (Browman and Goldstein, 1985). Articulatory studies typically use sections or windows that partition the movement curve on the basis of displacement minima and maxima. We compared such a windowing technique to another plausible windowing technique to see how the particular choice of window type might affect the values of the parameters.

This paper, then, reports on investigations into the above topics in which a computer program (PARFIT) was used to identify the values of the parameters in a mass–spring system. The mass–spring equation is

$$m\ddot{x} + b\dot{x} + k(x - x_0) = 0, \tag{1}$$

where m = mass, b = damping, k = stiffness, and x_0 = rest position. (In the work reported here, mass is assumed to be equal to 1.) The parameters extracted from the movement data correspond to the coefficients of the trigonometric form of the solution, shown in (2) below, which can be related analytically to the mass–spring equation above:

$$x(t) = e^{\alpha t}(A \cos \beta t + B \sin \beta t) + \text{d.c.} \tag{2a}$$

$$= \sqrt{A^2 + B^2} e^{\alpha t} \cos(\beta t - \theta) + \text{d.c.}, \tag{2b}$$

where α = growth, β = observed frequency, d.c. is the dc offset or constant level, A and B are a function of two selected values from the data, and θ is determined by A and B. The parameters in Eqs. (1) and (2) are related as follows:

$$\alpha = \frac{-b}{2}, \quad \beta = \frac{\sqrt{4k - b^2}}{2}, \quad \text{and d.c.} = x_0 \tag{3}$$

when mass = 1. Observed frequency (β) is related to natural frequency (ω_0) and damping ratio (d.rat.), itself a function of damping (b) and natural frequency:

$$\omega_0 = \sqrt{\frac{k}{m}}, \quad \beta = \omega_0 \sqrt{1 - (\text{d.rat})^2},$$

$$\text{where d.rat.} = \frac{b}{2\omega_0}. \tag{4}$$

One implication of the normalized mass in the equation of motion is that the stiffness is not being derived directly, but rather the frequency.

I. ANALYSES

The analyses were aimed first at determining the effect that linguistic factors such as stress, syllable position, and vowel quality have on the PARFIT parametrization of articulatory gestures in terms of natural frequency and damping ratio, and second at determining whether the effect of the linguistic factors remained stable across different damping ratios. In addition, the effect of window type on the results was investigated.

A. Procedure

PARFIT fits curves using a multidimensional Newton's method with a least-squared error criterion. That is, the PARFIT computer program attempts to find a set of parameter values such that the sum of squares of the differences between the positions generated by those values and the corresponding position data points is minimized. The parameters used were those in Eq. (2b). Further details about the algorithm used for fitting and tests of its validity can be found in McGowan *et al.* (1990); further details about the analyses of simulated data used for testing PARFIT can be found in Smith *et al.* (1991). For the simulated data, it was determined that stable results would be reliably attained only in the situation in which no more than two parameters were fit simultaneously. Therefore, analyses were run holding damping ratio constant, and allowing only the parameters of frequency and constant level to vary. Moreover, it was determined that different boundary condition options should be used for damping ratios between 0.0 and 0.8, and for damping ratios between 0.9 and 1.0.

The movement curves were divided into windows in two different ways to compare possible effects of different methods of sectioning curves (see Fig. 1). In the customary peak-to-valley size unit, referred to here as "peak windows," each window (portion of the curve) extended from the midpoint of a peak or valley (i.e., displacement maximum or minimum) to the midpoint of the next valley or peak. An alternative to this method, referred to as "CV windows," was also used, in order to include the relatively flat plateau regions around displacement extrema with the regions of movement between these plateaus. In the form ultimately used, each CV window extended from the right edge of a plateau region around a peak (or valley) to the right edge of the next valley (or peak). In the simulations, the plateaus began or ended within 1% of the range of amplitude after the extreme peak or valley of the movement curve.

Within each data file, each of the six windows was fit independently. In order to find the best fit, multiple fits of each window were performed (typically 11 fits for each window type, using fit damping ratios of 0.0 through 1.0 in increments of 0.1). Among fits of the same window made using different damping ratios, the hypothesis was that the one with the smallest squared error (referred to as "least error")

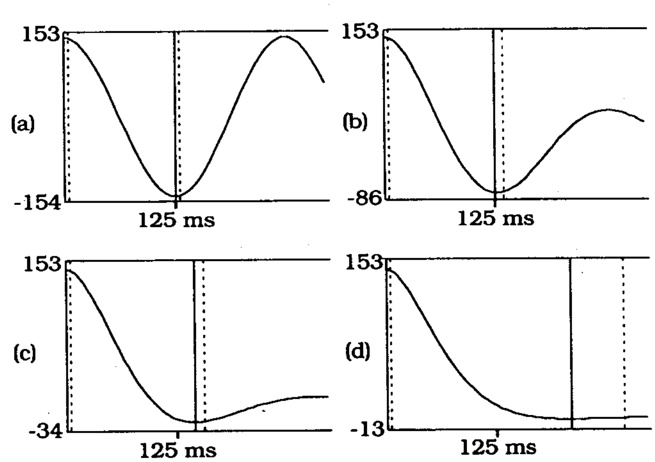

FIG. 1. Simulations at four damping ratios, (a) undamped, (b) 0.2, (c) 0.5, and (d) 0.8, with peak (solid lines) and CV (dotted lines) windows marked. The values on the vertical axis give the range of the display; the maximum (or minimum) in each figure is 10 machine units larger (or smaller) than the extreme value of the curve. The total duration displayed in each figure is 300 ms.

should provide the most accurate fit. This hypothesis was supported by the tests of the simulated data, in which it was determined that for both window types (peak and CV), the least error criterion provided estimates of damping ratio to within 0.1 of the correct value and mean estimates of natural frequency within 10% of the correct value (when used in conjunction with the appropriate boundary condition).

Selected fits were compared qualitatively using graphical representations to see how they diverged from the simulated curve. An example of such a comparison is shown in Fig. 2, which illustrates the results of using increasing damping ratio to fit an undamped curve (which has damping ratio = 0.0). At damping ratios 0.2 and 0.5, the fits shown in Fig. 2(a) and (b) diverge slightly from the simulated data curve, but they are quite close and in both cases fit the simulation (which consisted of half a cycle of an undamped sinusoidal curve) with approximately half of one cycle of the fit curve. As is apparent from Fig. 2(c), the close fit breaks down at fit damping ratio 0.8. In this case, the fit curve diverges drastically from the simulation, using much more than half of one cycle of the fit curve to fit the half-cycle simulated data window. Inaccuracy of the type shown by this fit was eliminated by using the least error criterion to select fits. Other kinds of inaccuracies were eliminated by placing constraints on the ranges of possible values for the parameters (e.g., observed frequency was constrained to be between 0 and 20 Hz; exponential growth was limited to between -200 and 5 s^{-1}, allowing only damped curves, with an allowance for a small positive amount of growth; and the constant level could not exceed 1.1 times the amplitude range of the window).

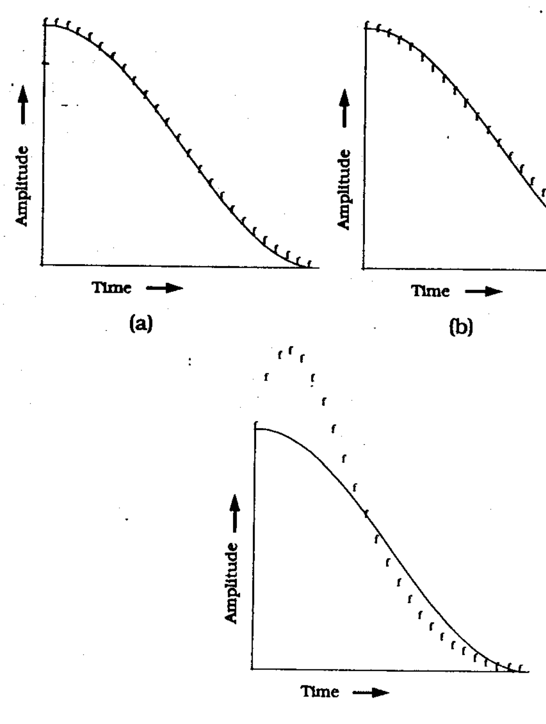

FIG. 2. The solid line represents a half-cycle of an undamped, simulated data curve (damping ratio = 0.0). The f's show fits of this curve using boundary condition 1 at damping ratios of (a) 0.2, (b) 0.5, and (c) 0.8.

In addition to the least error fits, sets of fits with fixed damping ratio were analyzed. Three different fixed damping ratios were used to sample the range between undamped (0.0) and critically damped (1.0): 0.2, 0.5, and 0.8. Analyses of variance using BMDP 4V were run on the extracted values for the damping ratios (least error fits only) and for frequency (for least-error fits and the three fixed-damping-ratio fits). Effects with p values below 0.05 were considered significant. Where main effects and interactions were both significant, tests of simple main effects were run to determine the extent to which the significance of the main effects held up in all conditions. In certain cases, post-hoc Newman–Keuls tests were also used.

B. Articulatory data

The articulatory movements used in the analyses were the vertical movements of the lower lip in space. One female speaker of American English was recorded, using a Selspot system with LEDs on the nose, upper lip, lower lip, and chin. The speaker produced four utterances containing contrasts among the linguistic factors discussed above: ['bibəbib], ['babəbab], [bibə'bib], and [babə'bab]. The utterances were produced in the carrier phrase "It's a — again." Eleven tokens of each were collected, except for the first utterance, for which 14 tokens were collected. The data were recorded on an FM tape recorder, then sampled at 200 Hz. The movement curves were smoothed using a 25-ms triangular window. (The tokens analyzed here were the same tokens as those analyzed in Browman and Goldstein, 1985.)

For each token, the data curve representing the vertical movement of the lower lip (i.e., lip plus jaw) was partitioned into windows corresponding to the opening and closing gestures, where for the present purposes the term "gesture" simply means a portion of a movement curve. Each window was marked in two ways, peak and CV, as described in the procedure. The edges of the CV windows were marked at points on the movement trace where the displacement from a movement extremum exceeded 0.58 mm. A sample utterance with the six windows marked is shown in Fig. 3(a), peak windows, and Fig. 3(b), CV windows. Figure 3(a) and (b) also shows how the gestures were associated with different values of the linguistic factors used in the statistical analyses, for peak and CV windows, respectively. The factors used in the statistical analyses were the following: syllable position in the utterance (1, 2, or 3), direction of movement (opening or closing), stress (stressed or unstressed), quality of the full vowel in the utterance (/i/ or /a/), and window type (peak or CV). The first four factors were considered to be linguistic factors, while window type was not.

In each utterance, stress fell on either the first full vowel (the first syllable) or the second full vowel (in the final syllable). The medial syllable was always reduced. If stress fell on the first vowel, gestures 1 through 3 (the opening into the first vowel through the opening into the schwa) were categorized as stressed, and gestures 4 through 6 (the closing out of the schwa through the closing out of the second full vowel) as unstressed. If stress was on the second full vowel, the first group of gestures was considered unstressed and the second stressed. Thus gestures for the medial schwa were catego-

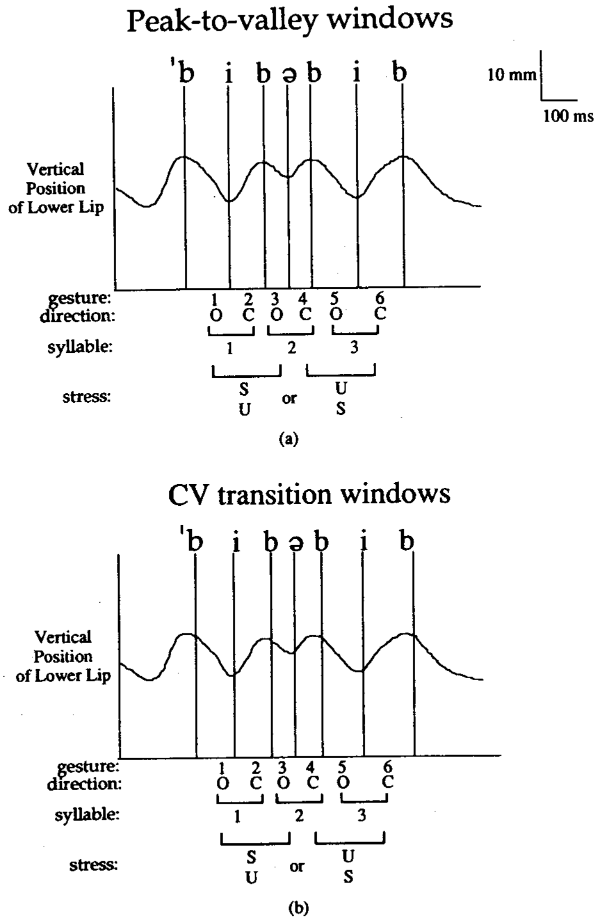

FIG. 3. A sample token of ['bibəbib] with the assignment of the values of the linguistic factors shown below: (a) peak windows, (b) CV windows. The direction of the gestures is either O (opening) or C (closing). Stress is indicated by S (stressed) and U (unstressed).

rized as stressed or unstressed depending on the neighboring syllable. This grouping, while not immediately intuitive, was chosen because results from preliminary analyses of these data had shown that the stress effects for the two gestures for the schwa tended to pattern with their adjacent full vowel, not with each other. In a similar fashion, all gestures in an utterance were categorized on the basis of the full vowel in the utterance (either /i/ or /a/), even though the gestures for the schwa were, of course, not directly involved in the production of the full vowel.

C. Results

Several preliminary analyses were performed to test the plausibility of the parametric analyses. First, measurements of the amplitude of the movements were made using peak windows; these amplitudes provided no surprises. Stressed gestures had larger amplitude movements than unstressed (mean of 7.53 vs 5.45 mm) and full vowels had larger amplitude than schwa (8.52 vs 2.31 mm). The stressed gesture with the largest mean amplitude was the opening into the first full vowel (11.72 mm), while the gesture with the smallest amplitude was the opening into the schwa when it followed an unstressed vowel (2.24 mm). Second, "phase angle" (the number of degrees of the curve that were fit, defined in terms of its natural frequency) was investigated to check whether the fit curve was a reasonable portion of the underlying curve. In general, the relation between natural frequency and phase angle in the least error fit was that expected (see Smith *et al.*, 1991). Moreover, the average phase angle was 187 deg, which means that the movements were being analyzed as being approximately half a cycle, as intended. Table I lists the phase angles in the least error fits by syllable position, direction, and window type. There was a strong tendency for larger phase angles to co-occur with higher frequencies. This was true in particular for stress, direction, and window type, all of which had significant effects on phase angle.

The parametric (PARFIT) analyses appeared to be well-behaved. Results of these analyses will be reported in detail below; here we will briefly sketch some general characteristics of the analyses. Looking first at the fits with fixed damping ratio (0.2, 0.5, and 0.8), it can be seen in Table II that as the fixed fit damping ratio increased, so did the mean values for natural frequency returned by the program. Looking next at the least error fits, which were selected from these and other fits made using a fixed damping ratio, as described in the procedure, the least error fits appeared to be quite accurate for these measured articulatory data. Although the size of error of the least error fits varied among the gestures, the mean errors for the fits across all gestures were of the same order of magnitude. For example, after normalizing the amplitude of the data to be between -1 and 1, the gesture with the smallest error was found to be the closing out of the first full vowel (syllable 1), with mean squared error of 0.00012, while the largest was the opening into the first full vowel (syllable 1), with a mean of 0.00059. Overall, opening gestures had 0.00040 mean-squared error, and closing gestures 0.00019.

Damping ratio (*least error fits only*). Recall that a single data file, which always consisted of six windows, could be fit with as many as six different damping ratios in the least error fits. The mean damping ratio for all the least error fits was 0.13, with values ranging from 0.0 to 0.57 across the categories defined by the factors. However, the linguistic factors did not have systematic effects on the damping ratios. Although the main effects of syllable, direction, stress, and vowel were significant overall, there were also many significant interactions. Simple main effects analyses (see Smith *et al.*, 1991) indicated that in fact all of the main effects were significant only in limited environments. The F values, degrees of freedom, and significance levels for the main analysis of variance of damping ratio are listed in Table III.

Natural frequency. The extracted values for natural frequency ranged from 1.84 Hz (for a stressed opening gesture

TABLE I. Phase angles.

Syllable	CV windows Closing	CV windows Opening	Peak windows Closing	Peak windows Opening
1	238	191	181	145
2	182	208	202	143
3	226	186	187	157

TABLE II. Mean frequencies at different damping ratios.

	Damping ratio	Mean frequency (in Hz)
(least error)	mean = 0.13	6.03
(fixed)	0.2	5.85
(fixed)	0.5	6.67
(fixed)	0.8	7.10

into the final full vowel, in syllable 3) to 12.56 Hz (for an unstressed opening gesture into the schwa, in syllable 2) in the least error fits. For natural frequency, the main effects of syllable, direction, stress, and window type were significant overall, but each interacted with other factors. The main effect of vowel did not reach significance. The *F* values, degrees of freedom, and significance levels are given in the Table IV; the simple main effects analyses are summarized in Table V. Selected simple main effects will also be examined individually below. In general, the results exhibited the same basic patterns among the frequency values for the fits at every damping ratio (whether least error or fixed at 0.2, 0.5, or 0.8), with the statistical significance of the factors somewhat reduced as the fit damping ratio increased, particularly for the 0.8 fits.

Simple main effects showed that the effect of *syllable* position on natural frequency was significant everywhere for the least error, 0.2, and 0.5 fits, and post-hoc Newman–Keuls tests showed that the natural frequency for each syllable was significantly different from that of each other syllable in each of these fits. These values are plotted in Fig. 4.

Simple main effects tests showed that the significance of *direction* held only in the syllables with full vowels (syllables 1 and 3) for least error, 0.2, and 0.5 fits. In these syllables, the frequency of closing gestures was significantly higher than that of opening gestures, as can be seen in Fig. 5(a)–(c) for the least error, 0.2, and 0.5 fits, respectively (the direction effect for the 0.8 fits, displayed in Fig. 5(d), was not statistically significant). The interaction of *syllable×direction*, which corresponds to the effect of individual gestures, was significant everywhere for these same three fits. Post-hoc Newman–Keuls tests for those fits in which the interaction of syllable×direction was significant showed that the individual gestures could be grouped into three to five significantly different groups, as indicated by the letters in Fig. 5(a)–(c). The lowest frequency gestures, labeled group A, were the opening gestures for the two full vowels. Group B, with frequency significantly higher than group A, is the closing gesture of the final syllable. The remaining gestures, generally associated with the schwa, had frequency values significantly higher than groups A and B. (Different letter names in this last group indicate further significant groupings.)

Simple main effects tests showed that the effect of *stress* on natural frequency was significant only in the first two syllables for the least error, 0.2, and 0.5 fits. In Fig. 6 it can be seen for least error fits that all stressed syllables had lower frequency values than the corresponding unstressed syllables, but in the final syllable this difference was so small that it was not significant statistically.

TABLE III. *F* values, degrees of freedom, and significance for analysis of variance of damping ratio for fits with least error (variable damping ratios). All two- and three-way interactions are shown; only the single four-way interaction that reached significance is shown. *** indicates significance of $p<0.001$, ** of $p<0.01$, and * of $p<0.05$.

	df	Least error
Syllable	2,516	47.77***
Direction	1,516	210.21***
Stress	1,516	7.72**
Vowel	1,516	13.13***
Window type	1,516	0.58
Syllable×direction	2,516	33.51***
Syllable×stress	2,516	...
Syllable×vowel	2,516	...
Syllable×window type	2,516	52.03***
Direction×stress	1,516	12.68***
Direction×vowel	1,516	4.27*
Direction×window type	1,516	14.92***
Stress×vowel	1,516	...
Stress×window type	1,516	22.88***
Vowel×window type	1,516	...
Syllable×direction×stress	2,516	...
Syllable×direction×vowel	2,516	...
Syllable×direction×window type	2,516	59.18***
Syllable×stress×vowel	2,516	...
Syllable×stress×window type	2,516	4.19*
Syllable×vowel×window type	2,516	...
Direction×stress×vowel	1,516	...
Direction×stress×window type	1,516	17.57***
Direction×vowel×window type	1,516	...
Stress×vowel×window type	1,516	8.82**
Syllable×direction×vowel×window type	2,516	6.60**

The effect of *window type* was found, by using simple main effects tests, to be significant only in opening gestures for the least error and 0.2 fits; it was significant everywhere in the 0.5 fits, and only in limited environments for the 0.8 fits. Higher frequency values were generally obtained using CV windows.

1. Anomalous gesture

The closing gesture out of the reduced syllable (syllable 2) often behaved anomalously, and the parameters extracted for that gesture were especially affected by the window type used for analysis. We will present the data for this gesture in some detail as a basis for the argument in the discussion that this apparently anomalous behavior in fact illuminates some important aspects of the analyses.

Damping ratio behaved differently with different window types for this gesture. This can be seen in the effect of direction on damping ratio, which was significant overall for CV windows but only in syllable 2 for peak windows. Figure 7 shows in general that the relation of opening and closing gestures is different for peak and CV windows, and in particular that the damping ratio for the closing gesture in the reduced middle syllable for peak windows is very different from the damping ratio of the other gestures. Post-hoc New-

TABLE IV. F values, degrees of freedom, and significance for analyses of variance of natural frequency of fits with least error (variable damping ratios), and of fits at three fixed damping ratios. All two- and three-way interactions are shown; only those four-way interactions that reached significance in at least one analysis are shown. *** indicates significance of $p<0.001$, ** of $p<0.01$, and * of $p<0.05$.

			Fit damping ratio		
	df	Least error	0.2	0.5	0.8
Syllable	2,516	132.58***	99.60***	205.55***	···
Direction	1,516	353.20***	106.54***	162.87***	24.10***
Stress	1,516	163.08***	73.85***	110.05***	21.22***
Vowel	1,516	···	···	···	···
Window type	1,516	159.60***	182.12***	452.06***	63.93***
Syllable × direction (= gesture)	2,516	115.60***	108.49***	95.61***	17.57***
Syllable × stress	2,516	26.28***	13.55***	21.86***	···
Syllable × vowel	2,516	···	···	5.11**	···
Syllable × window type	2,516	···	6.43**	6.70**	9.53***
Direction × stress	1,516	4.88*	···	···	···
Direction × vowel	1,516	···	4.04*	6.37*	···
Direction × window type	1,516	125.16***	145.85***	···	···
Stress × vowel	1,516	7.40**	···	11.70***	···
Stress × window type	1,516	7.25**	···	6.09*	···
Vowel × window type	1,516	···	···	···	···
Syllable × direction × stress	2,516	6.09**	···	···	···
Syllable × direction × vowel	2,516	16.06***	5.32**	8.86***	···
Syllable × direction × window type	2,516	64.56***	55.26***	39.99***	55.12***
Syllable × stress × vowel	2,516	···	···	3.02*	···
Syllable × stress × window type	2,516	···	···	···	4.09*
Syllable × vowel × window type	2,516	···	···	···	···
Direction × stress × vowel	1,516	···	···	4.51*	···
Direction × stress × window type	1,516	···	14.03***	···	···
Direction × vowel × window type	1,516	···	···	···	···
Stress × vowel × window type	1,516	···	···	···	11.49***
Syllable × direction × stress × vowel	2,516	···	···	3.90*	···
Syllable × direction × stress × window type	2,516	···	···	···	6.46**
Syllable × direction × stress × vowel × window type	2,516	···	···	···	4.82**

man–Keuls tests showed that, in peak windows, the damping ratio of the closing gesture of the middle syllable was significantly higher than the damping ratio of any other gesture; the other gestures did not differ significantly from one another.

Some of the patterns of natural frequency found among individual gestures were also changed for this anomalous gesture by the window type used in the analysis. For example, recall that the overall pattern was for opening gestures to have lower frequencies than the closing gestures within the same syllable. This effect was changed by window type in syllable 2 for least error, 0.2, and 0.8 fit damping ratios, as seen in Fig. 8(a), (b), and (d) (the gesture was anomalous in the 0.5 fits seen in Fig. 8(c) only for the CV opening/closing pattern). When analyzed with peak windows, as expected the frequency of the opening gesture was lower than that of the closing gesture. However, the reverse was true when this reduced syllable was analyzed with CV windows. The closing gesture for the reduced syllable was also unique in having a higher frequency with peak windows than with CV windows (for least error, 0.2, and 0.8 fit damping ratios).

D. Discussion

In this discussion, some of the results of the preceding analyses will be further examined in an attempt to determine whether the observed effects are the consequence of patterns in the data or of the particular analysis technique.

The effects of the linguistic factors on the natural frequency were remarkably robust across the various damping ratio assumptions (least error or fixed at 0.2, 0.5, 0.8), with only a decrease in the stability of the statistical significance when the damping ratio was increased to 0.8, a value very different from the average damping ratio of the data (0.13) as determined by the least error fits. That is, it appears that using an inappropriate (fixed) damping ratio did not alter the patterns of natural frequency values, but rather made some of them harder to discern (as indicated, for example, in Fig. 5, where the relationship between opening and closing gestures in a syllable is the same for all the damping ratios, although direction of gesture was not significant for the damping ratio of 0.8). The patterns that held at all the damping ratios included the following, none of them surprising:

TABLE V. Summary of significance of simple main effects for natural frequency, as shown by analyses of variance. The difference between frequency values was in the direction specified, unless a reversal is indicated in the chart by †. Limitations in the extent of significance are listed in the appropriate row. In some analyses the existence of multiple interactions made it necessary to break down the scope of the main effect's significance in more than one way, e.g., stress by syllable and by direction. "Gesture" is identical to the interaction syllable × direction, which identifies individual gestures in the utterance. (√ = significant in the environment named at left; -- = not significant; blank = interactions not significant; † = direction of effect reversed.)

	Least error	0.2	0.5[a]	0.8
Syllable (3 < 1 < 2)	√	√	√	--
Direction (open < close)				
Syllable 1	√	√	√	see below
Syllable 2	--	/i/†	/i/†	see below
Syllable 3	√	√	√	see below
Peak windows	√	√		Syll 2 × Unstr; Syll 3 × Str
CV windows	√	--		/i/; /a/ × Unstr; /a/ × Str × Syll 1
Gesture (see text)	√	√	√	
Stressed				CV
Unstressed				√
Stress (stressed < unstressed)				√
Syllable 1	√	√	√	
Syllable 2	√	√	√	
Syllable 3	--	--	--	
Open	√		/a/	
Close	√		√	
Vowel	--	--	--	--
Window type (peak < CV)			√	
Open	√	√		Syll 2
Close	--	--		Syll 1 × /i/; Syll 1 × Unstr; Syll 2 × /i/ × Unstr†; Syll 3 × /i/ × Str

[a]The interaction syllable × direction × stress × vowel ($F = 3.90$, $p < 0.05$) was not taken into account in determining the significances for the 0.5 fits.

Natural frequency was lower for stressed gestures than for unstressed gestures, even when the damping ratio of the fit (held fixed) was very far away from the least error damping ratio. In general, natural frequency was lower for opening than for closing gestures, even when the damping ratio was

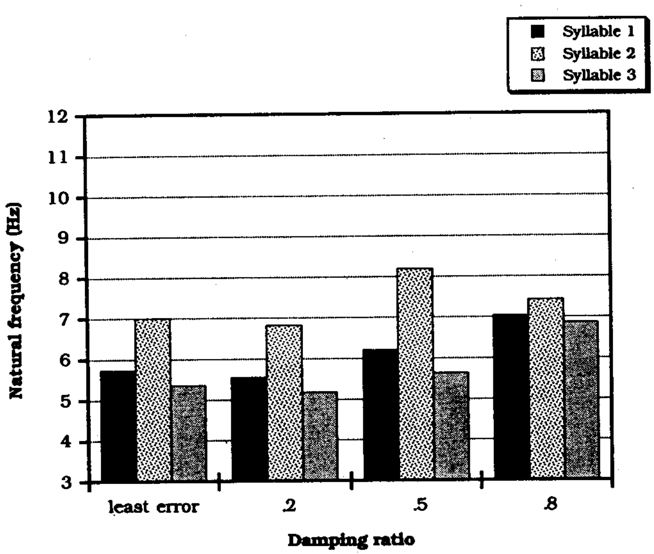

FIG. 4. Natural frequency values for each syllable in the least error fits and in fits at fixed damping ratios of 0.2, 0.5, and 0.8 for articulatory data.

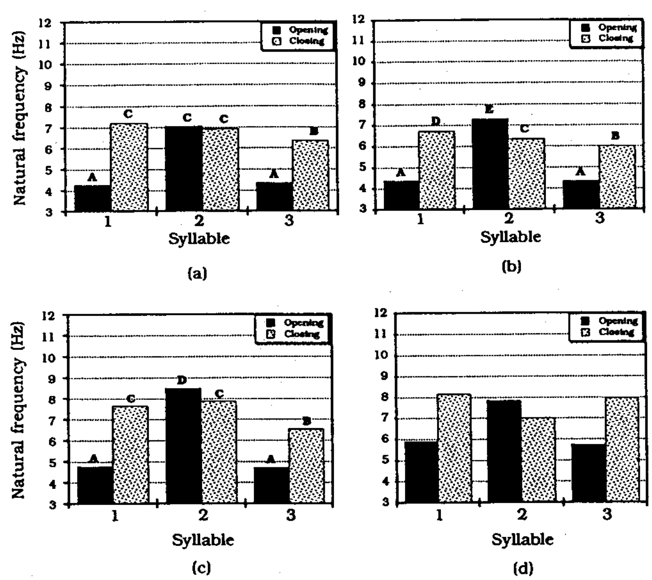

FIG. 5. Natural frequency values for opening (dark bars) and closing (dotted bars) gestures in each syllable for articulatory data. The letters indicate gestures that could be grouped by their significant differences (using Newman–Keuls tests); since the analyses were run separately for each of the four types of fits, identical letters are indicative of identical group membership only within each of the four subfigures. See text for limitations of significance: (a) least error fits; (b) fits with damping ratio 0.2; (c) fits with damping ratio 0.5; (d) fits with damping ratio 0.8.