

Unifying Amplitude and Phase Analysis: A Compositional Data Approach to Functional Multivariate Mixed-Effects Modeling of Mandarin Chinese

P. Z. HADJIPANTELIS, J. A. D. ASTON, H. G. MÜLLER, and J. P. EVANS

Mandarin Chinese is characterized by being a tonal language; the pitch (or F_0) of its utterances carries considerable linguistic information. However, speech samples from different individuals are subject to changes in amplitude and phase, which must be accounted for in any analysis that attempts to provide a linguistically meaningful description of the language. A joint model for amplitude, phase, and duration is presented, which combines elements from functional data analysis, compositional data analysis, and linear mixed effects models. By decomposing functions via a functional principal component analysis, and connecting registration functions to compositional data analysis, a joint multivariate mixed effect model can be formulated, which gives insights into the relationship between the different modes of variation as well as their dependence on linguistic and nonlinguistic covariates. The model is applied to the COSPRO-1 dataset, a comprehensive database of spoken Taiwanese Mandarin, containing approximately 50,000 phonetically diverse sample F_0 contours (syllables), and reveals that phonetic information is jointly carried by both amplitude and phase variation. Supplementary materials for this article are available online.

KEY WORDS: Functional data analysis; Linguistics; Multivariate linear mixed models; Phonetic analysis; Registration.

1. INTRODUCTION

Mandarin Chinese is one of the world's major languages (Central Intelligence Agency 2012) and is spoken as a first language by approximately 900 million people, with considerably more being able to understand it as a secondary language. Spoken Mandarin Chinese, in contrast to most European languages, is a tonal language (Su and Wang 2005). The modulation of the pitch of the sound is an integral part of the lexical identity of a word. Thus, any statistical approach of Mandarin pitch attempting to provide a pitch typology of the language, must incorporate the dynamic nature of the pitch contours into the analysis (Gu, Hirose, and Fujisaki 2006; Prom-On, Xu, and Thipakorn 2009).

Pitch contours, and individual human utterances generally, contain variations in both the amplitude and phase of the response, due to effects such as speaker physiology and semantic context. Therefore, to understand the speech synthesis process and analyze the influence that linguistic (e.g., context) and non-linguistic effects (e.g., speaker) have, we need to account for

variations of both types. Traditionally, in many phonetic analyses, pitch curves have been linearly time normalized, removing effects such as speaker speed or vowel length, and these time normalized curves are subsequently analyzed as if they were the original data (Xu and Wang 2001; Aston, Chiou, and Evans 2010). However, this has a major drawback: potentially interesting information contained in the phase is discarded as pitch patterns are treated as purely amplitude variational phenomena.

In a philosophically similar way to Kneip and Ramsay (2008), we model both phase and amplitude information jointly and propose a framework for phonetic analysis based on functional data analysis (FDA; Ramsay and Silverman 2005) and multivariate linear mixed-effects (LME) models (Laird and Ware 1982). Using a single multivariate model that concurrently models amplitude, phase, and duration, we are able to provide a phonetic typology of the language in terms of a large number of possible linguistic and nonlinguistic effects, giving rise to estimates that conform directly to observed data. We focus on the dynamics of F_0 ; F_0 is the major component of what a human listener identifies as speaker pitch (Jurafsky and Martin 2009) and relates to how fast the vocal folds of the speaker vibrate. We use two interlinked sets of curves; one set consisting of time normalized F_0 amplitude curves and a second set containing their corresponding time-registration/warping functions registering the original curves to a universal time-scale. Using methodological results from the compositional data literature (Aitchison 1982), a principal component analysis of the centered log ratio of the time-registration functions is performed. The principal component scores from the amplitude curves and the time warping functions along with the duration of the syllable are then jointly modeled through a multivariate LME framework.

© P. Z. Hadjipantelis, J. A. D. Aston, H. G. Müller, and J. P. Evans

This is an Open Access article. Non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly attributed, cited, and is not altered, transformed, or built upon in any way, is permitted. The moral rights of the named author(s) have been asserted.

P. Z. Hadjipantelis is PhD, Centre for Research in Statistical Methodology; Centre for Complexity Science, University of Warwick, Coventry CV4 7AL, UK (E-mail: p.z.hadjipantelis@warwick.ac.uk). J. A. D. Aston is Professor, Centre for Research in Statistical Methodology, University of Warwick, Coventry CV4 7AL, UK and Department of Pure Maths and Mathematical Statistics, Statistical Laboratory, University of Cambridge, Cambridge, UK (E-mail: j.aston@statslab.cam.ac.uk). H. G. Müller is Professor, Department of Statistics, University of California, Davis, CA 95616 (E-mail: hgmuller@ucdavis.edu). J. P. Evans is PhD, Institute of Linguistics, Academia Sinica, Taipei, Taiwan (E-mail: jonathan@gate.sinica.edu.tw). JADA's research was supported by the Engineering and Physical Sciences Research Council (EP/K021672/2). HGM's research was supported by NSF grants DMS-1104426 and DMS-1228369. JPE's research was supported by National Science Council (Taiwan) grant NSC 100-2628-H-001-008-MY4.

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/r/jasa.

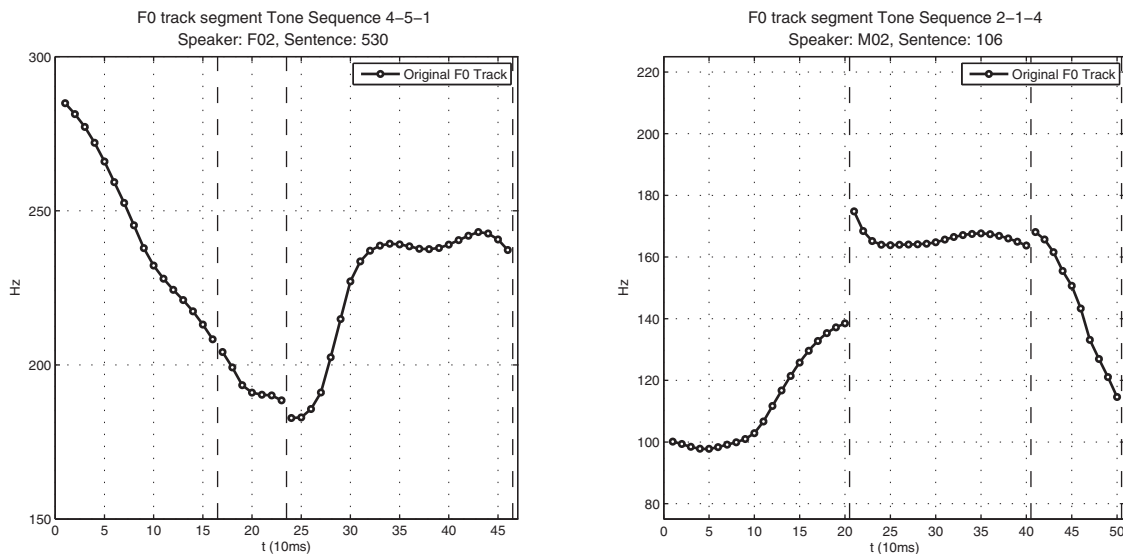


Figure 1. An example of triplet trajectories from speakers *F02* and *M02* over natural time. *F*(emale)02 tonal sequence: 4-5-1, *M*(ale)02 tonal sequence: 2-1-4; Mandarin Chinese rhyme sequences [oŋ-@-iou] and [ien-in-l], respectively. See supplementary material for full contextual covariate information.

One aspect of note in our modeling approach is that it is based on a compositional representation of the warping functions. This representation is motivated by viewing the registration functions on normalized time domains as cumulative distribution functions, with derivatives that are density functions, which in turn can be approximated by histograms arbitrarily closely in the L^2 norm. We may then take advantage of the well-known connection between histograms and compositional data (Leonard 1973; Pawlowsky-Glahn and Egozcue 2006).

The proposed model is applied to a large linguistic corpus of Mandarin Chinese consisting of approximately 50,000 individual syllables in a wide variety of linguistic and nonlinguistic contexts. Due to the large number of curves, computational considerations are of critical importance to the analysis. The dataset is prohibitively large to analyze with usual multilevel computational implementations (Hadfield 2010; Bates et al. 2013), so a specific computational approach for the analysis of large multivariate LME models is developed. Using the proposed model, we are able to identify a joint model for Mandarin Chinese that serves as a typography for spoken Mandarin. This study thus provides a robust and flexible statistical framework describing intonation properties of the language.

The article proceeds as follows. In Section 2, a short review of the linguistic properties of Mandarin will be given. General statistical methodology for the joint modeling of phase and amplitude functions will then be outlined in Section 3, including its relation to compositional data analysis and other methods of modeling phase and amplitude. Section 4 contains the analysis of the Mandarin corpus where it will be seen that the model not only provides a method for determining the role of linguistic covariates in the synthesis of Mandarin, but also allows comparisons between the estimated and the observed curves. Finally, the last section contains a short discussion of the future prospects of FDA in linguistics. Further details of the analysis implemented are given in the supplementary material.

2. PHONETIC ANALYSIS OF MANDARIN CHINESE

We focus our attention on modeling fundamental frequency (F_0) curves. The amplitude of F_0 , usually measured in Hz, quantifies the rate/frequency of the speaker's vocal folds' vibration and is an objective measure of how high or low the speaker's voice is. In this study, the observation units of investigation are brief syllables: F_0 segments that typically span between 120 and 210 msec (Figure 1) and are assumed to be smooth and continuous throughout their trajectories. Linguistically our modeling approach of F_0 curves is motivated by the intonation model proposed by Fujisaki (2004) where linguistic, paralinguistic, and nonlinguistic features are assumed to affect speaker F_0 contours. Another motivation for our rationale of combining phase and amplitude variation comes from the successful usage of hidden Markov models (HMM; Rabiner 1989; Yoshioka et al. 2012) in speech recognition and synthesis modeling. However, unlike the HMM approach, we aim to maintain a linear modeling framework favored by linguists for its explanatory value (Baayen, Davidson, and Bates 2008; Evans et al. 2010) and suitability for statistical modeling.

Usual approaches segment the analysis of acoustic data. First one applies a "standard" dynamic time warping (DTW) treatment to the sample using templates (Sakoe 1979), registers the data in this new universal time scale, and then continues with the analysis of the variational patterns in the synchronized speech utterances (Latsch and Netto 2011). In contrast, we apply functional principal component analysis (FPCA; Castro, Lawton, and Sylvestre 1986) to the "warped" F_0 curves and also to their corresponding warping functions, the latter being produced during the curve registration step. These functional principal component scores then serve as input for using a multivariate LME model, allowing a joint modeling of both the phase and amplitude variations.

Table 1. Covariates examined in relation to F_0 production in Taiwanese Mandarin. Tone variables in a 5-point scale representing tonal characterization, 5 indicating a toneless syllable, with 0 representing the fact that no rhyme precedes the current one (such as at the sentence start). Reference tone trajectories are shown in the supplementary material section: Linguistic Covariate Information

Effects	Values	Meaning	Notation mark
<i>Fixed effects</i>			
Previous tone	0:5	Tone of previous syllable, 0 no previous tone present	tn_{previous}
Current tone	1:5	Tone of syllable	tn_{current}
Following tone	0:5	Tone of following syllable, 0 no following tone present	tn_{next}
Previous consonant	0:3	0 is voiceless, 1 is voiced, 2 not present, 3 sil/short pause	cn_{previous}
Next consonant	0:3	0 is voiceless, 1 is voiced, 2 not present, 3 sil/short pause	cn_{next}
B2	linear	Position of the B2 index break in sentence	$B2$
B3	linear	Position of the B3 index break in sentence	$B3$
B4	linear	Position of the B4 index break in sentence	$B4$
B5	linear	Position of the B5 index break in sentence	$B5$
Sex	0:1	1 for male, 0 for female	Sex
Duration	linear	10 sec of msec	Duration
Rhyme type	1:37	Rhyme of syllable	rhyme _{<i>t</i>}
<i>Random effects</i>			
Speaker	$N(0, \sigma_{\text{speaker}}^2)$	Speaker effect	SpkrID
Sentence	$N(0, \sigma_{\text{sentence}}^2)$	Sentence effect	Sentence

We analyze a comprehensive speech corpus of Mandarin Chinese. The Sinica Continuous Speech Prosody Corpora (COSPRO; Tseng, Cheng, and Chang 2005) was collected at the Phonetics Lab of the Institute of Linguistics in Academia Sinica and consists of nine sets of speech corpora. We focus our attention on the COSPRO-1 corpus; the phonetically balanced speech database. COSPRO-1 was designed to specifically include all possible syllable combinations in Mandarin based on the most frequently used 2- to 4-syllable lexical words. Additionally, it incorporates all the possible tonal combinations and concatenations. It therefore offers a high quality speech corpus that, in theory at least, encapsulates all the prosodic effects that might be of acoustic interest. Specifically, we analyze 54,707 fully annotated “raw” syllabic F_0 curves that were uttered by a total of five native Taiwanese Mandarin speakers (two males and three females). Each speaker uttered the same 598 predetermined sentences having a median length of 20 syllables; each syllable had on average 16 readings.

In total, aside from Speaker and Sentence information, associated with each F_0 curve are covariates of break index (within word (B2), intermediate (B3), intonational (B4), and utterance (B5) segments), its adjacent consonants, its tone, and rhyme type (Table 1). In our work, all of these variables serve as potential scalar covariates and with the exception of break counts, the fixed covariates are of categorical form. The break (or pause) counts, representing the number of syllables between successive breaks of a particular type, are initialized at the beginning of the sentence and are subsequently reset every time a corresponding or higher order break occurs. They represent the perceived degree of disjunction between any two words, as defined in the ToBi annotations (Beckman, Hirschberg, and Shattuck-Hufnagel 2006). Break counts are very significant as physiologically a break has a resetting effect on the vocal folds’ vibrations; a qualitative description of the break counts is provided in Table S.1 of the supplementary material.

3. STATISTICAL METHODOLOGY

3.1 A Joint Model

The application of functional data analysis (FDA) in the field of Phonetics, while not wide-spread, is not unprecedented; previous functional data analyses included lip-motion (Ramsay et al. 1996), analysis of prosodic effects (Lee, Byrd, and Krivokapic 2006), speech production (Koenig, Lucero, and Perlman 2008) as well as basic language investigation based solely on amplitude analysis (Aston, Chiou, and Evans 2010). FDA is, by design, well suited as a modeling framework for phonetic samples as F_0 curves are expected to be smooth. Concurrent phase and amplitude variation is expected in linguistic data and as phonetic datasets feature “dense” measurements with high signal-to-noise ratios (Ramsay and Silverman 2005), FDA naturally emerges as a statistical framework for F_0 modeling. Nevertheless in all phonetic studies mentioned above, the focus of the phonetic analysis has been almost exclusively the amplitude variations (the size of the features on a function’s trajectory) rather than the phase variation (the location of the features on a function’s trajectory) or the interplay between the two domains.

To alleviate the limitation of only considering amplitude, we use the formulation presented by Tang and Müller (2009) and introduce two types of functions, w_i and h_i associated with our observed curve y_i , $i = 1, \dots, N$, where y_i is the i th curve in the sample of N curves. For a given F_0 curve y_i , w_i is the amplitude variation function on the domain $[0, 1]$ while h_i is the monotonically increasing phase variation function on the domain $[0, 1]$, such that $h_i(0) = 0$ and $h_i(1) = 1$. For generic random phase variation or warping functions h and time domains $[0, T]$, T also being random, we consider time transformations $u = h^{-1}(\frac{t}{T})$ from $[0, T]$ to $[0, 1]$ with inverse transformations $t = Th(u)$. Then, the measured curve y_i over the interval $t \in [0, T_i]$ is assumed to be of the

form

$$\begin{aligned} y_i(t) &= w_i \left(h_i^{-1} \left(\frac{t}{T_i} \right) \right) \Leftrightarrow \\ w_i(u) &= y_i(T_i h_i(u)), \quad i = 1, \dots, N, \end{aligned} \tag{3.1}$$

where $u \in [0, 1]$ and T_i is the duration of the i th curve. A curve y_i is viewed as a realization of the amplitude variation function w_i evaluated over u , with the mapping $h_i^{-1}(\cdot)$ transforming the scaled real time t onto the universal/sample-wide time-scale u . In addition, each curve can depend on a set of covariates, fixed effects X_i , such as the tone being said, and random effects Z_i , where such random effects correspond to additional speaker and context characteristics. While each individual curve has its own length T_i , which is directly observed, the lengths entering the functional data analysis are normalized and the T_i are subsequently included in the modeling as part of the multivariate linear mixed effects framework.

In our application, the curves y_i are associated with various covariates, for example, tone, speaker, and sentence position. These are incorporated into the model via the principal component scores that result from adopting a common principal component approach (Flury 1984; Benko, Härdle, and Kneip 2009), where we assume common principal components (across covariates) for the amplitude functions and another common set (across covariates) for phase functions (but these two sets can differ). We use a common PCA framework with common mean and eigenfunctions so that all the variation in both phase and amplitude is reflected in the respective FPC scores. These ideas have been previously used in a regression setting (although not in the context of registration) (Benko, Härdle, and Kneip 2009; Aston, Chiou, and Evans 2010; Chen and Müller 2014). As will be discussed in Section 4.2, this is not a strong assumption in this application. Of the covariates likely present in the model, tone is known to affect the shape of the curves (indeed it is in the phonetic textual representation of the syllable), and therefore the identification of warping functions is carried out within tone classes as opposed to across the classes, as otherwise very strong (artificial) warpings will be introduced.

As a direct consequence of our generative model (Equation 1), w_i dictates the size of a given feature and h_i^{-1} dictates the location of that feature for a particular curve i . We assume that w_i and h_i are both elements of $L^2[0, 1]$. The w_i can be expressed in terms of a basis expansion:

$$w_i(u) = \mu^w(u) + \sum_{j=1}^{\infty} A_{i,j}^w \phi_j(u), \tag{3.2}$$

where $\mu^w(u) = E\{w(u)\}$, ϕ_j is the j th basis function, and $A_{i,j}^w$ is the coefficient for the i th amplitude curve associated with the j th basis function. The h_i are a sample of random distribution functions that are square integrable but are not naturally representable in a basis expansion in the Hilbert space $L^2[0, 1]$, since the space of distribution functions is not closed under linear operations. A common approach to circumvent this difficulty is to observe that $\log(\frac{d}{du}h(u))$ is not restricted and can be modeled as a basis expansion in $L^2[0, 1]$. A restriction however is that the densities h_i have to integrate to 1, therefore the functions $s_i(u) = \log(\frac{d}{du}h_i(u))$ are modeled with the unrestricted basis

expansion:

$$s_i(u) = \mu^s(u) + \sum_{j=1}^{\infty} A_{i,j}^s \psi_j(u), \tag{3.3}$$

where $\mu^s(u)$, ψ_k , and $A_{i,j}^s$ are defined analogously to $\mu^w(u)$, ϕ_k , and $A_{i,j}^w$, respectively, but for the warping rather than amplitude functions. A transformation step is then introduced to satisfy the integration condition, which yields the representation:

$$h_i(u) = \frac{\int_0^u e^{s_i(u')} du'}{\int_0^1 e^{s_i(u')} du'} \tag{3.4}$$

for the warping function h_i . Clearly different choices of bases will give rise to different coefficients A , which then can be used for further analysis. A number of different parametric basis functions can be used as basis; for example, Grabe et al. advocated the use of Legendre polynomials (Grabe, Kochanski, and Coleman 2007) for the modeling of amplitude. We advocate the use of a principal component basis for both w_i and s_i in Equations 2.3 and 3.3, as will be discussed in the next sections, although any basis can be used in the generic framework detailed here. However, a principal components basis does provide the most parsimonious basis in terms of a residual sum of squares like criterion (Ramsay and Silverman 2005). We note that to ensure statistical identifiability of the model (Equation 3.1), several regularity assumptions were introduced in Tang and Müller (2008, 2009), including the exclusion of essentially flat amplitude functions w_i for which time-warping cannot be reasonably identified, and more importantly, assuming that the time-variation component that is reflected by the random variation in h_i and s_i asymptotically dominates the total variation. In practical terms, we will always obtain well-defined estimates for the component representations in Equations 3.2 and 3.3, and their usefulness hinges critically on their interpretability; see Section 5.

For our statistical analysis, we explicitly assume that each covariate X_i influences, to different degrees, all of the syllable's components/modes as well as influencing the syllable's duration T_i . Additionally, as mentioned above in accordance with the Fujisaki model, we assume that each syllable component includes speaker-specific and sentence-specific variational patterns; we incorporate this information in the covariates Z_i . Then the general form of our model for a given sample curve y_i of duration T_i with two sets of scalar covariates X_i and Z_i is

$$E\{w_i(u)|X_i, Z_i\} = \mu^w(u) + \sum_{j=1}^{\infty} E\{A_{i,j}^w|X_i, Z_i\}\phi_j(u), \tag{3.5}$$

and

$$E\{s_i(u)|X_i, Z_i\} = \mu^s(u) + \sum_{j=1}^{\infty} E\{A_{i,j}^s|X_i, Z_i\}\psi_j(u). \tag{3.6}$$

Truncating expansions (3.5) and (3.6), to M_w and M_s components, respectively, is a computational necessity and simplifies the implementation. Curves are then reduced to finitely many scores $A_{i,j}^w$, $A_{i,j}^s$ and these score vectors then act as surrogate data for curves w_i and s_i . The final joint model for amplitude, phase,

and syllable duration is then formulated as

$$E\{[A_i^w, A_i^s, T_i]|X_i, Z_i\} = X_i B + Z_i \Gamma, \quad \Gamma \sim N(0, \Sigma_\Gamma), \quad (3.7)$$

where A_i^w and A_i^s are the vectors of component coefficients for the i th sample. Here, B (a $k \times p$ matrix, where k is the number of fixed effects in the model and p is the number of multivariate components in the mixed effects model, $p = M_w + M_s + 1$, where the “1” arises from the additional duration component in the model) is the parameter matrix of the fixed effects and Γ (a $l \times p$ matrix, where l is the number of random effects in the model) contains the coefficients of the random effects and is assumed to have mean zero, while Σ_Γ is the covariance matrix of the amplitude, phase, and duration components with respect to the random effects.

The process as a whole is summarized in [Figure 2](#).

3.2 Amplitude Modeling

In our study, amplitude analysis is conducted through a functional principal component analysis of the amplitude variation functions. Qualitatively, the random amplitude functions w_i are the time-registered versions of the original F_0 samples. Using FPCA, we determine the eigenfunctions that correspond to the principal modes of amplitude variation in the sample and then use a finite number of eigenfunctions corresponding to the largest eigenvalues as a truncated basis, so that representations in this basis will explain a large fraction of the total variation. Specifically, we define the kernel C^w of the covariance operator as

$$C^w(u, u^*) = E\{(w(u) - \mu^w(u))(w(u^*) - \mu^w(u^*))\} \quad (3.8)$$

and by Mercer’s theorem (Mercer 1909), the spectral decomposition of the symmetric amplitude covariance function C^w can be written as

$$C^w(u, u^*) = \sum_{p_w=1}^{\infty} \lambda_{p_w} \phi_{p_w}(u) \phi_{p_w}(u^*), \quad (3.9)$$

where the eigenvalues λ_{p_w} are ordered by declining size and the corresponding eigenfunction is ϕ_{p_w} . Additionally, the eigenvalues λ_{p_w} allow the determination of the total percentage of variation exhibited by the sample along the p th principal component and whether the examined component is relevant for further analysis. As will be seen later, the choice of the number of components is based on acoustic criteria (Black and Hunt 1996; Kochanski 2006) with direct interpretation for the data, such that components that are not audible are not considered. Having fixed M_w as the number of ϕ modes/eigenfunctions to retain, we use ϕ to compute A_{i,p_w}^w , the amplitude projections scores associated with the i th sample and its p_w th corresponding component (Equation (3.10)) as

$$A_{i,p_w}^w = \int \{w_i(u) - \mu^w(u)\} \phi_{p_w}(u) dt, \quad (3.10)$$

where as before $\mu^w(u) = E\{w(u)\}$,

where a suitable numerical approximation to the integral is used for practical analysis.

3.3 Phase Modeling

When examining the warping functions it is important to note that we expect the mean of the random warping function to correspond to the identity (i.e., the case of no warping). Therefore, assuming their domains are all normalized to $[0,1]$, with $t = Th(u)$, this assumption is

$$u = E\{h(u)\}, \quad (3.11)$$

and this allows one to interpret deviations of h from the identity function as phase distortion. This clearly also applies conceptually when working with the function $s(u)$. As with the amplitude analysis, phase analysis is carried out using a principal component analysis approach. Using the eigenfunctions of the random warping functions s_i , we identify the principal modes of variation of the sample and use those modes as a basis to project our data to a finite subspace. Directly analogous to the decomposition of C^w , the spectral decomposition of the phase covariance function C^s is

$$C^s(u, u^*) = \sum_{p_s=1}^{\infty} \lambda_{p_s} \psi_{p_s}(u) \psi_{p_s}(u^*), \quad (3.12)$$

where the ψ_{p_s} are the eigenfunctions and the λ_{p_s} are the eigenvalues, ordered in declining order. The variance decomposition through eigencomponents is analogous to that for the amplitude functions. As before we will base our selection processes not on an arbitrary threshold based on percentages but on an acoustic perceptual criteria (Quene 2007; Jacewicz, Fox, and Wei 2010) for perceivable speed changes. If one retains M_s eigencomponents in the expansion (3.6), the corresponding functional principal component scores for the i th warping or phase variation function w_i are obtained as (3.13):

$$A_{i,p_s}^s = \int \{s_i(u) - \mu^s(u)\} \psi_{p_s}(u) dt, \quad (3.13)$$

where as before $\mu^s(u) = E\{s(u)\}$.

It is worth stressing that our choice of the number of components to retain will be dictated by an external criterion that is specific to the phonetic application, rather than being determined by a purely statistical criterion such as fraction of total variance explained. Purely data-driven approaches have been developed (Minka 2001) as well as a number of different heuristics (Cangelosi and Goriely 2007) for less structured applications, where no natural and interpretable choice is available.

3.4 Sample Time-Registration

The estimation of the phase variation/warping functions is as in Tang and Müller (2008), as implemented in the routine WFPCA in PACE (Tang and Müller 2009). There, one defines the pairwise warping function $g_{i',i}(t) = h_i(h_i^{-1}(t))$ as the 1-to-1 mapping from the i th curve’s time-scale to that of the i' th by minimizing a distance (usually selected as L^2 distance) by warping the time-scale of the i th curve as closely as possible to that of the i' th curve. The inverse of the average $g_{i',i}(\cdot)$ (Equation (3.15)) for a curve i then can be shown to yield a consistent estimate of the time-warping function h_i that is specific for curve y_i and corresponds to a map between individual-specific warped time to absolute time (Tang and Müller 2008).

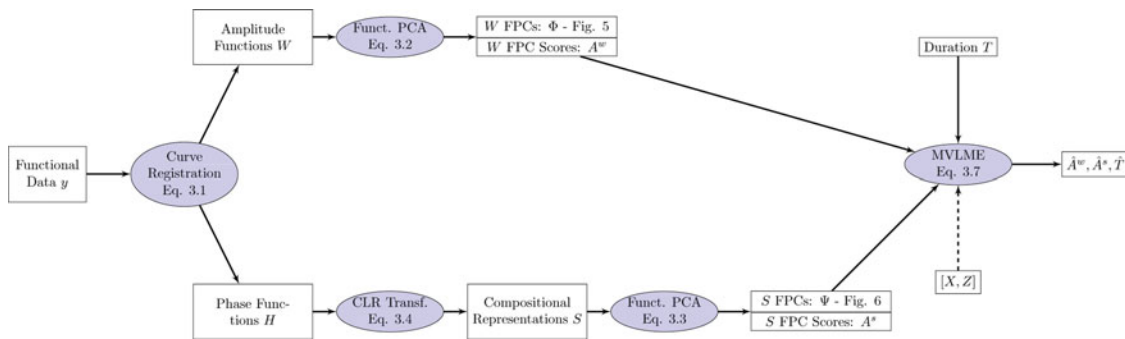


Figure 2. Summary of the overall estimation procedure resulting in the estimates of the functional principal components and scores via the covariates in the linear mixed effect model.

The $g_{i',i}(\cdot)$, being time-scale mappings, have a number of obvious restrictions on their structure. First, $g_{i',i}(0) = 0$ and $g_{i',i}(1) = 1$. Second, they should be monotonic, that is, $g_{i',i}(t_j) \leq g_{i',i}(t_{j+1}), 0 \leq t_j < t_{j+1} \leq 1$. Finally, $E[g_{i',i}(t)] = t$. This final condition is necessary for representing the warping function h_i through its inverse h_i^{-1} :

$$h_i^{-1}(t) = E[g_{i',i}(t)|h_i(t)] \tag{3.14}$$

with sample version:

$$\hat{h}_i^{-1}(t) = \frac{1}{N^*} \sum_{i'=1}^{N^*} \hat{g}_{i',i}(t), \quad N^* \leq N, \tag{3.15}$$

where N^* is the number of sample pairwise registrations used to obtain the estimate. In small datasets, all the curves can be used for the pairwise comparisons that lead to (3.15), but in a much larger dataset such as the one in our phonetic application, only a random subsample of curves of size N^* is used to obtain the estimates for computational reasons.

Aside from the pairwise alignment framework we employ (Tang and Müller 2008), we have identified at least two alternative approaches based on different metrics, the square-root velocity function metric (Kurtek et al. 2012) or area under the curve normalization metric (Zhang and Müller 2011), which can be used interchangeably, depending on the properties of the warping that are considered most important in specific application settings. Indeed it has been seen that considering warping and amplitude functions together, based on the square-root velocity metric, can be useful for classification problems (Tucker, Wu, and Srivastava 2013). However, we need to stress that each method makes some explicit assumptions to overcome the non-identifiability between the h_i and w_i (Equation (3.1)) and this can lead to significantly different final estimates.

3.5 Compositional Representation of Warping Functions

To apply the methods to obtain the time warping functions in Section 3.4 and their functional principal component representations in Section 3.3, we still need a suitable representation of individual warping functions that ensures that these functions have the same properties as distribution functions. For this purpose, we adopt step function approximations of the warping functions h_i , which are relatively simple yet can approximate any distribution function arbitrarily closely in the L^2 or sup norms by choosing the number of steps large enough. A natural

choice for the steps is the grid of the data recordings, as the phonetic data are available on a grid. The differences in levels between adjacent steps then give rise to a histogram that represents the discretized warping functions.

This is where a novel connection to the proposed compositional decompositions arises. Based on standard compositional data methodology (centered log-ratio transform) (Aitchison 1982), the first difference $\Delta h_{i,j} = h_i(t_{j+1}) - h_i(t_j)$ of a discretized instance of h_i over an $(m + 1)$ -dimensional grid is used to evaluate s_i as

$$s_{i,j} = \log \frac{\Delta h_{i,j}}{(\Delta h_{i,1} \cdot \Delta h_{i,2} \dots \Delta h_{i,m})^{\frac{1}{m}}} \quad j = \{1, \dots, m\} \tag{3.16}$$

the reverse transformation being:

$$h_{i,j+1} = \frac{e^{s_{i,j}}}{\sum_j e^{s_{i,j}}}, \quad h_{i,1} = 0. \tag{3.17}$$

This ensures that monotonicity ($h_{i,j} < h_{i,j+1}$), and boundary requirements ($h_{i,1} = 0, h_{i,m+1} = 1$) are fulfilled as required in the pairwise warping step; this compositional approach guarantees that an evaluation will always remain in the space of warping functions. The sum of the first differences of all discretized h_i warping functions is equal to a common C , and thus Δh_i is an instance of compositional data (Aitchison 1982).

We can then employ the centered log-ratio transform for the analysis of the compositional data, essentially dividing the components by their geometric means and then taking their logarithms. The centered log-ratio transform has been the established method of choice for the variational analysis of compositional data; alternative methods such as the additive log-ratio (Aitchison 1982) or the isometric log-ratio (Egozcue et al. 2003) are also popular choices. In particular, the centered log-ratio, as it sums the transformed components to zero by definition, presents itself as directly interpretable in terms of “time-distortion,” negative values reflecting deceleration and positive values acceleration in the relative phase dynamics. Clearly, this summation constraint imposes a certain degree of collinearity in our transformed sample (Filzmoser, Hron, and Reimann 2009); nevertheless it is the most popular choice of compositional data transformation (Aitchison 1983; Aitchison and Greenacre 2002) and allows direct interpretation as mentioned above.

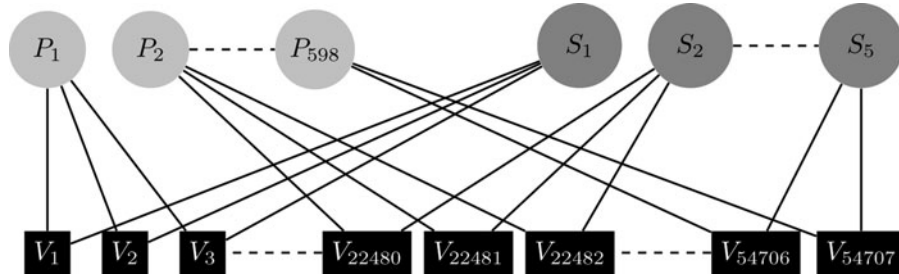


Figure 3. The multivariate mixed effects model presented exhibits a crossed (nonbalanced) random structure. The vowel-rhyme curves (V) examined are cross-classified by their linguistic (Sentence— P_i) and their nonlinguistic characterization (Speaker— S_i).

3.6 Further Details on Mixed Effects Modeling

Given an amplitude-variation function w_i , its corresponding phase-variation function s_i , and the original F_0 curve duration T_i , each sample curve is decomposed into two mean functions (one for amplitude and one for warping) and an $M_w + M_s + 1 := p$ vector space of partially dependent measurements arising from the scores associated with the eigenfunctions along with the duration. Here, M_w is the number of eigencomponents encapsulating amplitude variations, M_s is the number of eigencomponents carrying phase information, and the 1 refers to the curves’ duration. The final linear mixed effect model for a given sample curve y_i of duration T_i and sets of scalar covariates X_i and Z_i is

$$\begin{aligned} [A_{i,k}^w, A_{i,m}^s, T_i] &= X_i B + Z_i \Gamma + E_i, \\ \Gamma &\sim N(0, \Sigma_\Gamma), \quad E \sim N(0, \Sigma_E), \end{aligned} \quad (3.18)$$

Σ_E being the diagonal matrix of measurement error variances (Equation S.2). The covariance structures Σ_Γ and Σ_E are of particular forms; while Σ_E (Equation S.2) assumes independent measurement errors, the random effects covariance (Equation S.1) allows a more complex covariance pattern (see supplementary material for relevant equations).

As observed in previous work (Hadjipantelis, Aston, and Evans 2012), the functional principal components for the amplitude process are uncorrelated among themselves and so are those for the phase process. However, between phase and amplitude they are expected to be correlated, and will also be correlated with time T . Therefore, the choice of an unstructured covariance for the random effects is necessary; we have found no theoretical or empirical evidence to believe any particular structure such as a compound symmetric covariance structure, for example, is present within the eigenfunctions and/or duration. Nevertheless our framework would still be directly applicable if we choose another restricted covariance (e.g., compound symmetry) structure and if anything it will become computationally easier to investigate as the number of parameters would decrease.

Our sample curves are concurrently included in two nested structures: one based on “speaker” (nonlinguistic) and one based on “sentence” (linguistic) (Figure 3). We therefore have a crossed design with respect to the random-effects structure of the sample (Brumback and Rice 1998; Aston, Chiou, and Evans

2010), which suggests the inclusion of random effects:

$$A_{N \times p} = X_{N \times k} B_{k \times p} + Z_{N \times l} \Gamma_{l \times p} + E_{N \times p}, \quad (3.19)$$

where p is the multivariate dimension, k is the number of fixed effects, and l is the number of random effects, as before. This generalization allows the formulation of the conditional estimates as

$$A | \Gamma \sim N(XB + Z\Gamma, \Sigma_E) \quad (3.20)$$

or unconditionally and in vector form for \vec{A} as

$$\begin{aligned} \vec{A}_{Np \times 1} &\sim N((I_p \otimes X) \vec{B}_{Np \times 1}, \Lambda_{Np \times Np}), \\ \Lambda &= (I_p \otimes Z)(\Sigma_\Gamma \otimes I_l)(I_p \otimes Z)^T + (\Sigma_E \otimes I_N), \end{aligned} \quad (3.21)$$

where X is the matrix of fixed effects covariates, B is the matrix of fixed effects coefficients, Z is the matrix of random effects covariates, Γ is the matrix of random effects coefficients (a sample realization dictated by $N(0, \Sigma_\Gamma)$), $\Sigma_\Gamma = D_\Gamma^{\frac{1}{2}} P_\Gamma D_\Gamma^{\frac{1}{2}T}$, the random effects covariance matrix, D_Γ is the diagonal matrix holding the individual variances of random effects, P_Γ is the correlation matrix of the random effects between the series in columns i, j , and Σ_E is the diagonal measurement errors covariance matrix. Kronecker products (\otimes) are used to generate the full covariance matrix Λ of \vec{A} as the sum of the block covariance matrix for the random effects and the measurement errors.

3.7 Estimation

Estimation is required in two stages: obtaining the warping functions and multivariate mixed effects regression estimation. Requirements for the estimation of pairwise warping functions $g_{k,i}$ were discussed in Section 3.4. In practical terms, these requirements mean that: (i) $g_{k,i}(\cdot)$ needs to span the whole domain, (ii) we cannot go “back in time,” that is, the function must be monotonic, and (iii) the time-scale of the sample is the average time-scale followed by the sample curves. With these restrictions in place, we can empirically estimate the pairwise (not absolute) warping functions by targeting the minimizing time transformation function $g_{k,i}(\cdot)$ as $\hat{g}_{k,i}(t) = \operatorname{argmin}_g D(y_k, y_i, g)$, where the “discrepancy” cost function D is defined as

$$\begin{aligned} D_\lambda(y_k, y_i, g) &= E \left\{ \int_0^1 (y_k(g(t); T_k) - y_i(t; T_i))^2 \right. \\ &\quad \left. + \lambda(g(t) - t)^2 dt | y_k, y_i, T_k, T_i \right\}, \end{aligned} \quad (3.22)$$

λ being an empirically evaluated nonnegative regularization constant, chosen in a similar way to Tang and Müller (2008); see also Ramsay and Li (1998); T_i and T_k being used to normalize the curve lengths. Intuitively the optimal $g_{k,i}(\cdot)$ minimize the differences between the reference curve y_i and the “warped” version of y_k subject to the amount of time-scale distortion produced on the original time-scale t by $g_{k,i}(\cdot)$. Having a sufficiently large sample of N^* pairwise warping functions $g_{k,i}(\cdot)$ for a given reference curve y_i , the empirical internal time-scale for y_i is given by Equation (3.15), the global warping function h_i being easily obtainable by simple inversion of h_i^{-1} . It is worth noting that in Mandarin, each tone has its own distinct shape; their features are not similar and therefore should not be aligned. For this reason, the curves were warped separately per tone, that is, realizations of Tone1 curves were warped against other realizations of Tone1 only, the same being applied to all other four tones. In order for the minimization in (3.22) to be well defined, it is essential to have a finite-dimensional representation for the time transformation/warping functions g . Such a representation is provided by the compositional centered log transform and this makes it possible to implement the minimization.

Finally to estimate the mixed model via the model’s likelihood, we observe that usual maximum likelihood (ML) estimation underestimates the model’s variance components (Patterson and Thomson 1971). We therefore use restricted maximum likelihood (REML); this is essentially equivalent to taking the ML estimates for our mixed model after accounting for the fixed effects X . The restricted maximum (log)likelihood estimates are given by maximizing the following formula:

$$L_{REML}(\theta) = -\frac{1}{2}[p(N - k) \log(2\pi) + \log(|\Psi|) + \vec{\Omega}^T \Psi^{-1} \vec{\Omega}], \quad (3.23)$$

where $\Psi = K^T \Lambda K$ and $\Omega = K^T A$; with K being the “whitener” matrix such that $0 = K^T(I_p \otimes X)$ (Searle, Casella, and McCulloch 2006). Based on this, we concurrently estimate the random effect covariances while taking into account the possible nondiagonal correlation structure between them. Nevertheless because we “remove” the influence of the fixed effects to compare models with different fixed effects structures we would need to use ML rather REML estimates. Standard mixed-effects software such as lme4 (Bates et al. 2013), nlme (Pinheiro et al. 2013), and MCMCglmm (Hadfield 2010) either do not allow the kinds of restrictions on the random effects covariance structures that we require, as they are not designed to model multivariate mixed effects models, or computationally are not efficient enough to model a dataset of this size and complexity; we were therefore required to write our own implementation for the evaluation of REML/ML. Exact details about the optimization procedure used to do this are given in the supplementary material section: Computational aspects of multivariate mixed effects regression.

4. DATA ANALYSIS AND RESULTS

4.1 Sample Preprocessing

It is important as a first step to ensure F_0 curves are “smooth,” that is, they possess “one or more derivatives” (Ramsay and Sil-

verman 2005). In line with Chiou, Müller, and Wang (2003), we use a locally weighted least squares smoother to fit local linear polynomials to the data and produce smooth data-curves interpolated upon a common time-grid on a dimensionless interval $[0, 1]$. Guo (2002) presented a smoothing framework producing comparable results by employing smoothing splines. The form of the kernel smoother used is as in Chiou, Müller, and Wang (2003) with fixed parameter bandwidth estimated using cross-validation (Izenman 2008) and Gaussian kernel function.

The curves in the COSPRO sample have an average of 16 readings per case, hence the number of grid points chosen was 16. The smoother bandwidth was set to 5% of the relative curve length. As is common in a dataset of this size, occasional missing values have occurred and curves having 5% or more of the F_0 readings missing were excluded from further analysis. These missing values usually occurred at the beginning or the end of a syllable’s recording and are most probably due to the delayed start or premature stopping of the recording. During the smoothing procedure, we note each curve’s original time duration (T_i) so it can be used within the modeling. At this point, the F_0 curve sample is not yet time-registered but has been smoothed and interpolated to lie on a common grid.

4.2 Model Presentation and Fitting

As mentioned in Section 2, the data consisted of approximately 50,000 sample curves. However, as can be seen in Figure S.2, in noncontextual situations, the tones have simple and distinct shapes. Therefore, registration was not performed on the dataset in its entirety but rather using each tone class as its own registration set. This raises an interesting discussion as to whether the curves are now one common sample, or rather a group of five separate samples. However, if we assume that the five tone groups all have common means and principal components (Benko, Härdle, and Kneip 2009) for both amplitude and phase variations, then this alleviates any issues with the use of separate registrations. This assumption substantially simplifies the model and is not particularly restrictive in that the ability of the vocal folds to produce very different pitch contours is limited, and as such it is likely that common component contours are present in each group.

It is possible to develop many different linguistic models for this data. However, the following model is proposed, as it accounts for all the linguistic effects that might be present in a dataset of this form (Hadjipantelis, Aston, and Evans 2012), which is a particular case of (3.18) where the covariates are now specified:

$$[A_{i,k}^w, A_{i,m}^s, T_i] = \{[tn_{previous} * tn_{current} * tn_{next}] + [cn_{previous} * tn_{current} * cn_{next}] + [(B2) + (B2)^2 + (B2)^3 + (B3) + (B3)^2 + (B3)^3 + (B4) + (B4)^2 + (B4)^3 + (B5) + (B5)^2 + (B5)^3] * Sex + [rhyme_t]_i B + \{[Sentence] + [SpkrID]\}_i \Gamma + E_i. \quad (4.1)$$

Standard Wilkinson notation (Wilkinson and Rogers 1973) is used here for simplicity regarding the interaction effects; [K*L] represents a short-hand notation for [K + L + K:L] where the

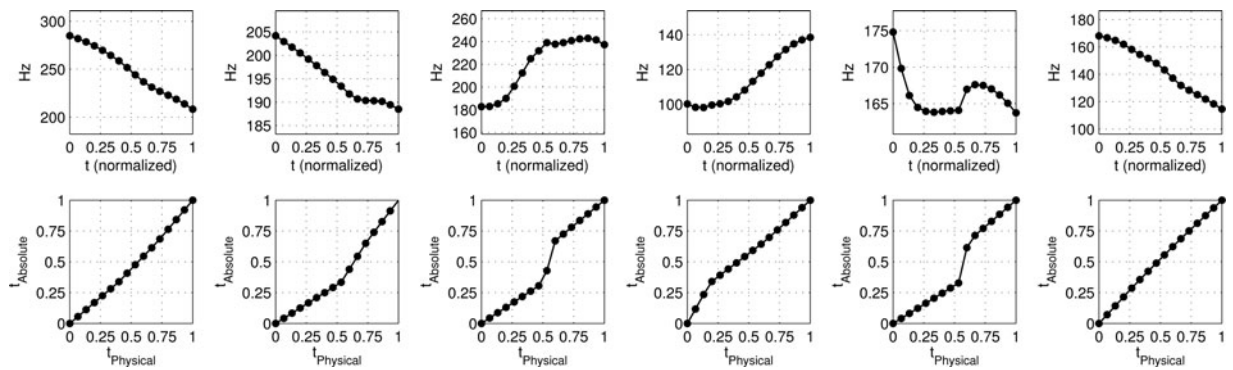


Figure 4. Corresponding amplitude variation functions w (top row) and phase variation functions h (bottom row) functions for the triples shown in Figure 1.

colon specifies the interaction of the covariates to its left and right (Baayen 2008). First, examining the fixed effects structure, we incorporate the presence of tone-triplets and of consonant:tone:consonant interactions. Both types of three-way interactions are known to be present in Mandarin Chinese and to significantly dictate tonal patterns (Xu 1999; Torgerson 2005). We also look at break counts, our only covariate that is not categorical. A break’s duration and strength significantly affects the shape of the F_0 contour and not just within a rhyme but also across phrases. Break counts are allowed to exhibit squared and cubic patterns as cubic downdrift has been previously observed in Mandarin studies (Aston, Chiou, and Evans 2010; Hadjipantelis, Aston, and Evans 2012). We also model breaks as interacting with the speaker’s sex as we want to provide the flexibility of having different curvature declination patterns among male and female speakers. This partially alleviates the need to incorporate a random slope as well as a random intercept in our mixed model’s random structure. The final fixed effect we examine is the type of rhyme uttered. Each syllable consists of an initial consonant or \emptyset followed by a rhyme. The rhyme contains a vowel followed by $-\emptyset/ -n/ -\eta$. The rhyme is the longer and more sonorous part of the syllable during which the

tone is audible. Rhyme types are the single most linguistically relevant predictors for the shape of F_0 ’s curve as when combined together they form words, with words carrying semantic meaning.

Examining the random effects structure we incorporate speaker and sentence. The inclusion of speaker as a random effect is justified as factors of age, health, neck physiology, and emotional condition affect a speaker’s utterance and are mostly immeasurable but still rather “subject-specific.” Additionally, we incorporate Sentence as a random effect since it is known that pitch variation is associated with the utterance context (e.g., commands have a different F_0 trajectory than questions). We need to note that we do not test for the statistical significance of our random effects; we assume they are “given” as any linguistically relevant model has to include them. Nevertheless if one wished to access the statistical relevance of their inclusion, the χ^2 mixtures framework used by Lindquist et al. (2012) provides an accessible approach to such a high-dimensional problem, as resampling approaches (bootstrapping) are computationally too expensive in a dataset of the size considered here. Fixed effects comparisons are more straightforward; assuming a given random-effects structure, Akaike information criterion

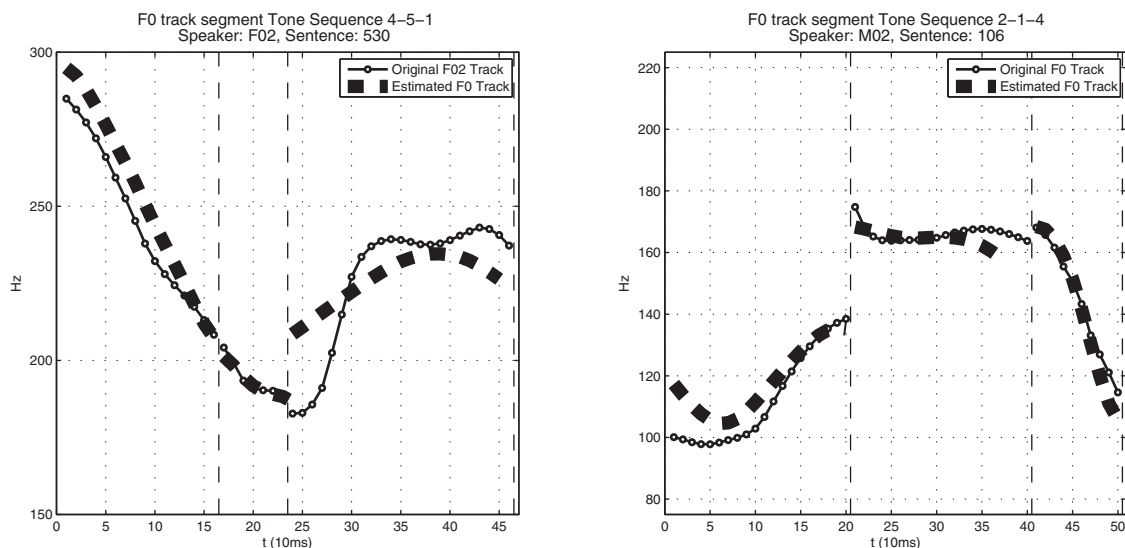


Figure 5. Functional estimates (continuous curves) are shown superimposed of the corresponding original discretized speaker data over the physical time domain t .

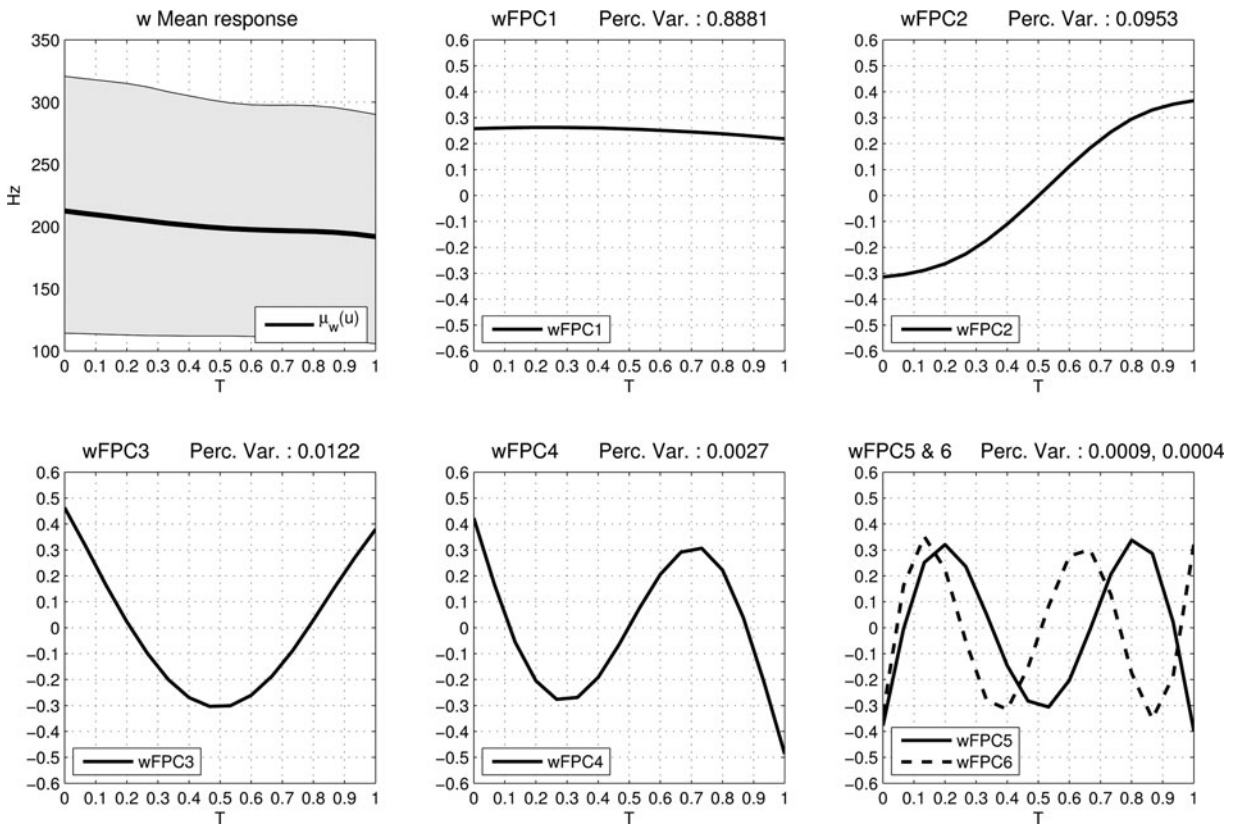


Figure 6. W (amplitude) eigenfunctions Φ : mean function ([0.05, 0.95] percentiles shown in gray) and first, second, third, fourth, fifth, and sixth functional principal components (FPCs) of amplitude.

(AIC)-based methodology can be directly applied (Greven and Kneib 2010). Fitting the models entails maximizing REML of the model (Equation (3.23)).

Our findings can be grouped into three main categories, those from the amplitude analysis, those from the phase, and those from the joint part of the model. Some examples of the curves produced by the curve registration step are given in Figure 4. However, overall, as can be seen in Figure 5, there is a good correspondence between the model estimates and the observed data when the complete modeling setup is considered. Small differences in the estimates can be ignored due to the just noticeable difference (JND) criteria (see below). The only noticeable departure between the estimates and the observed data is in the third segment of Figure 5 (left). The sinusoidal difference in the measured data that is not in the estimate can be directly attributed to the exclusion of amplitude PC's five and six, as these were below the JND criteria. The continuity difference in the observed curves is not enforced by the model and is hence not as prominent in the estimates. The general shape is the same but the continuity yields a sharper change in the observed data than is expected. It would be of great interest in future research to extend the ideas of registration to models where both amplitude and warping functions are included as bivariate dependent random functions.

Empirical findings from the amplitude FPCA: The first question one asks when applying any form of dimensionality reduction is how many dimensions to retain, or more specifically in the case of FPCA how many components to use. We take a perceptual approach. Instead of using an arbitrary percentage of

variation, we calculate the minimum variation in Hz each FPC can actually exhibit (Tables 2 and 3). Based on the notion of JND (Buser and Imbert 1992), we use for further analysis only eigenfunctions that reflect variation that is actually detectable by a standard speaker (F_0 JND: ≈ 10 Hz; $M_w = 4$). The empirical w FPCs (Figure 6) correspond morphologically to known Mandarin tonal structures (Figure S.2) increasing our confidence in the model. Looking into the analogy between components and reference tones with more details, w FPC1 corresponds closely to Tone 1, w FPC2 can be easily associated with the shape of Tones 2 and 4 and w FPC3 corresponds to the U -shaped structure shown in Tone 3. w FPC4 appears to exhibit a sinusoid pattern that can be justified as necessary to move between different tones in certain tonal configurations (Hadjipantelis, Aston, and Evans 2012).

Empirical findings from the phase FPCA: Again the first question is how many components to retain. Based on existing JND in tempo studies (Quene 2007; Jacewicz, Fox, and Wei 2010), we choose to follow their methodology for choosing the number of “relevant” components (tempo JND: $\approx 5\%$ relative distortion; $M_s = 4$). We focus on percentage changes on the transformed domain over the original phase domain as it is preferable to conduct principal component analysis (Aitchison 1983); s FPCs also corresponding to “standard patterns” (Figure 7). s FPC1 and s FPC2 exhibit a typical variation one would expect for slow starts and/or trailing syllable utterances where a decelerated start leads to an accelerated ending of the word—a catch-up effect—and vice versa. s FPC3 and s FPC4, on the other hand, show more complex variation patterns that

Table 2. Percentage of variances reflected from each respective FPC (first 9 shown). Cumulative variance in parenthesis

	Amplitude/(w)	Phase/(s)
FPC ₁	88.67 (88.67)	49.40 (49.40)
FPC ₂	10.16 (98.82)	19.25 (68.65)
FPC ₃	0.75 (99.57)	9.02 (77.68)
FPC ₄	0.22 (99.80)	6.53 (84.19)
FPC ₅	0.10 (99.90)	4.34 (88.53)
FPC ₆	0.05 (99.94)	2.98 (91.51)
FPC ₇	0.02 (99.97)	2.32 (93.83)
FPC ₈	0.01 (99.98)	1.96 (95.79)
FPC ₉	0.01 (99.99)	1.29 (97.08)

Table 3. Actual deviations in Hz from each respective FPC (first 9 shown). Cumulative deviance in parenthesis. (Human speech auditory sensitivity threshold ≈ 10 Hz)

	Amplitude/(w)
FPC ₁	121.16(121.16)
FPC ₂	66.52 (187.68)
FPC ₃	31.22 (218.90)
FPC ₄	17.50 (236.40)
FPC ₅	9.00 (245.39)
FPC ₆	4.86 (250.26)
FPC ₇	3.64 (253.90)
FPC ₈	2.71 (256.61)
FPC ₉	1.96 (258.56)

are most probably rhyme specific (e.g., /-ia/) or associated with uncommon sequences (e.g., silent pause followed by a Tone 3) and do not have an obvious universal interpretation. While the curves in Figure 7 are not particularly smooth due to the discretized nature of the modeling, as can be seen in Figure S.1 in the supplementary material, the resulting warping functions after transformation are smooth.

Empirical findings from the Multi-Variate Linear Mixed Effects (MVLME) analysis: The most important joint findings are the correlation patterns presented in the covariance structures of the random effects as well as their variance amplitudes. A striking phenomenon is the small, in comparison with the residual amplitude, amplitudes of the Sentence effects (Table 4). This

goes to show that pitch as a whole is much more speaker dependent than context dependent. It also emphasizes why certain pitch modeling algorithms focus on the simulations of “neck physiology”(Taylor 2000; Fujisaki 2004; Louw and Barnard 2004).

In addition to that we see some linguistically relevant correlation patterns in Figure 8 (also see S.25–S.26 in the supplementary material). For example, wFPC2 and duration are highly correlated both in the context of Speaker and Sentence-related variation. The shape of the second wFPC is mostly associated with linguistic properties (Hadjipantelis, Aston, and Evans 2012) and a syllable’s duration is a linguistically relevant

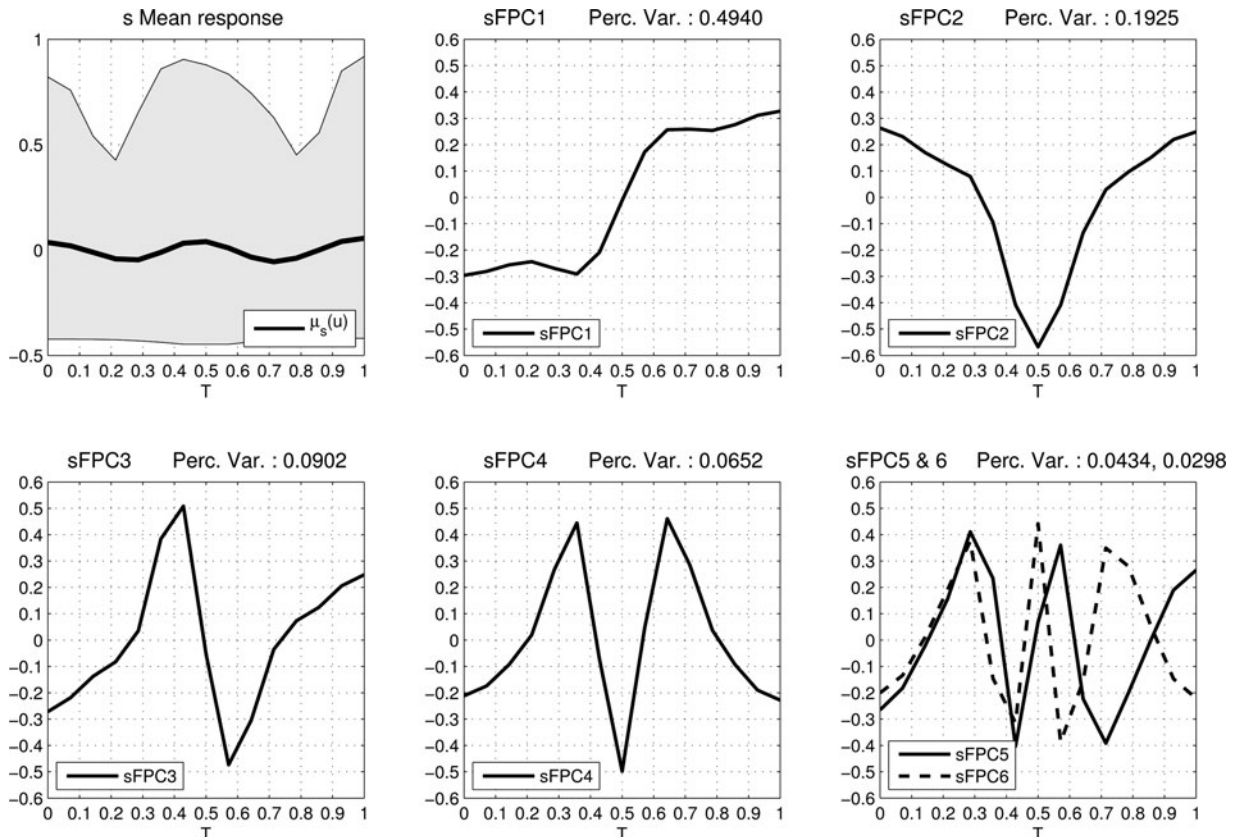


Figure 7. (Phase) eigenfunctions Ψ : mean function ([0.05, 0.95] percentiles shown in gray) and first, second, third, fourth, fifth, and sixth functional principal components (FPCs) of phase. Roughness is due to differentiation and finite grid; the corresponding warping functions in their original domain are given in Figure S.1 in the supplementary material.

Table 4. Random effects std. deviations

Estimate	wFPC ₁	wFPC ₂	wFPC ₃	wFPC ₄	Duration	sFPC ₁	sFPC ₂	sFPC ₃	sFPC ₄
Speaker	89.245	6.326	3.655	1.330	2.806	0.289	0.023	0.022	0.030
Sentence	38.674	4.059	0.045	0.102	0.043	0.049	0.043	0.042	0.043
Residual	114.062	44.386	15.399	10.072	4.481	0.959	0.591	0.431	0.370

property itself. As wFPC2 is mostly associated with the slope of syllable’s F_0 trajectory, it is unsurprising that changes in the slope affect the duration. Moreover, looking at the signs we see that while the Speaker influence is negative, in the case of Sentence, it is positive. That means that there is a balance on how variable the length of an utterance can be to remain comprehensible (so, e.g., when a speaker tends to talk more slowly than normal, the effect of the Sentence will be to “accelerate” the pronunciation of the words in this case). (See supplementary material for Σ_{R_i} ’s definitions.) In relation to that, in the speaker random effect, sFPC₁ is also correlated with duration as well as wFPC₂; yielding a triplet of associated variables. Looking specifically to another phase component, sFPC₂ indicating mid syllable acceleration or deceleration that allow for changes in the overall pitch patterns, is associated with a syllable’s duration, this being easily interpreted by the face that such changes are modulated by alterations in the duration of the syllable itself. Complementary to these phenomena is the relation between the syllable duration and wFPC₁ sentence-related variation. This correlation does not appear in the speaker effects and thus is likely due to more linguistic rather than physiological changes in the sample. As mentioned previously, wFPC1 can be thought of as dictating pitch-level placement, and the correlation implies that higher-pitched utterances tend to last longer. This is not contrary to the previous finding; higher F_0 placements are necessary for a speaker to utter a more pronounced slope differential and obviously need more time to be manifested.

Interestingly, a number of lower magnitude correlation effects appear to associate wFPC₁ and sFPCs. This is something

that needs careful interpretation. wFPC₁ is essentially “flat” (Figure 6, upper middle panel), and as such cannot be easily interpreted when combined with registration functions. Nevertheless this shows the value in our joint modeling approach for these data. We concurrently account for all these correlations during model estimation and, as such, our estimates are less influenced by artifacts in individual univariate FPCs.

Examining the influence of fixed effects, the presence of adjacent consonants was an important feature for almost every component in the model. Additionally, certain “domain-specific” fixed effects emerged also. The syllable’s rhyme type appeared to significantly affect duration; the break-point information to influence the amplitude of the F_0 curve and specific consonant-vowel-consonant (C-V-C) triplets to play a major role for phase. Phase also appeared to be related to the rhyme types but to a lesser extent (table of \hat{B} and associated standard errors available in <https://tinyurl.com/COSPRO-Betas>).

More specifically regarding duration of the F_0 curve, certain rhyme types (e.g., /-oŋ/, /-iEn/) gave prominent elongation effects while others (e.g., /-u/, /-f/) were associated with shorter curves. These are high vowels, meaning that the jaw is more closed and the tongue is nearer to the top of the mouth than for low vowels. It is to be expected that some rhymes are shorter than others and that high vowels with no following nasal consonant would indeed be the shorter ones. The same pattern of variability in the duration was associated with the adjacent consonants information; when a vowel was followed by a consonant, the F_0 curve was usually longer while when the consonant preceded a vowel, the F_0 curve was shorter. Amplitude-related components are significantly affected by the utterances’ break-

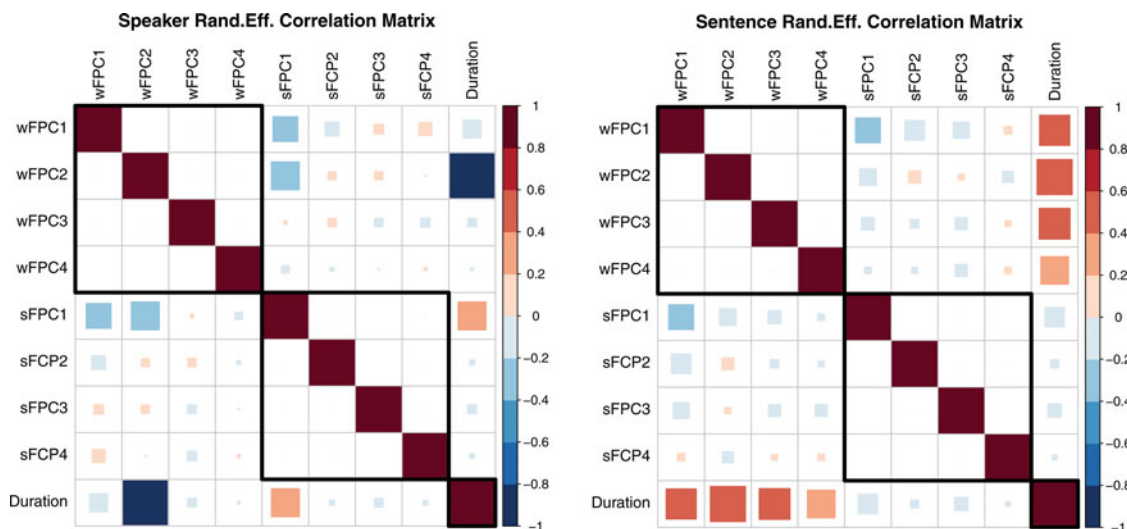


Figure 8. Random effects correlation matrices. The estimated correlation between the variables of the original multivariate model (Equation (3.19)) is calculated by rescaling the variance-covariance submatrices Σ_{R_1} and $\Sigma_{R_2}^2$ of Σ_{Γ} to unit variances. Each cell i, j shows the correlation between the variance of component in row i that of column j ; row/columns 1–4: wFPC1–4, row/columns 5–8: sFPC1–4, row/columns 9: Duration.

type information; particularly B2 and B3 break types. This is not a surprising finding; a pitch trajectory, to exhibit the well-established presence of “down-drift” effects (Fujisaki 2004), needs to be associated with such variables. As in the case of duration, the presence of adjacent consonants affects the amplitude dynamics. Irrespective of its type (voiced or unvoiced), the presence of consonant before or after a rhyme led to an “overall lowering” of the F_0 trajectory. Tone type and the sex of the speaker also influenced the dynamics of amplitude but to a lesser degree. Finally examining phase, it is interesting that most phase variation was mainly due to the adjacent consonants and the rhyme type of the syllable; these also being the covariates affecting duration. This confirms the intuition that as both duration and phase reflect temporal information, they would likely be affected by the same covariates. More specifically a short or a silent pause at the edge of rhyme caused that edge to appear decelerated while the presence of a consonant caused that edge to be accelerated. As before, certain rhymes (e.g., /-a/, /-ai/) gave more pronounced deceleration–acceleration effects. Tone types, while very important in the case of univariate models for amplitude (Hadjipantelis, Aston, and Evans 2012), did not appear significant in this analysis individually; they were usually significant when examined as interactions. However, this again illustrates the importance of considering joint models versus marginal models, as it allows a more comprehensive understanding of the nature of covariate effects.

In addition, we have reimplemented the main part of the analysis using the area under the curve methodology of Zhang and Müller (2011) that had previously been considered in Liu and Müller (2004) (results shown in supplementary material, Section S.6) and while the registration functions obtained are different, the analysis resulted in almost identical insights for the linguistic roles of w_i and s_i , again emphasizing the need to consider a joint model.

5. DISCUSSION

Linguistically our work establishes the fact that when trying to give a description of a language’s pitch, one needs to take care of amplitude and phase covariance patterns while correcting for linguistic (Sentence) and nonlinguistic (Speaker) effects. This need was prominently demonstrated by the strong correlation patterns observed (Figure 8). Clearly we do not have independent components in our model and therefore a joint model is appropriate. This has an obvious theoretical advantage in comparison to standard linguistic modeling approaches such as MO-MEL (Hirst and Espesser 1993) or the Fujisaki model (Mixdorff 2000; Fujisaki 2004) where despite the use of splines to model amplitude variation, phase variation is ignored.

Focusing on the interpretation of our results, it is evident that the covariance between phase and amplitude is mostly due to nonlinguistic (Speaker-related) rather than linguistic features (Sentence-related). This is also reflected in the dynamics of duration, where the influence is also greater (than the Sentence-related). Our work as a whole presents a first coherent statistical analysis of pitch incorporating phase, duration, and amplitude modeling into a single overall approach.

One major statistical issue with the interpretation of our results is due to the inherent identifiability problem of curve reg-

istration. It is not possible, without extra assumptions, to determine two functions (amplitude and phase) from one sampled curve. While this is a problem in general, especially for the relatively simply structured pitch functions that we consider here, nonidentifiability of the decomposition of total variation into warping and amplitude variation is a well-known issue. This is in contrast with situations where functions have distinct structures such as well-defined peaks (Kneip and Gasser 1992). In any case, identifiability usually needs to be enforced by model assumptions or algorithmically. We use pairwise registration, for which identifiability conditions have been given in Tang and Müller (2008). In practice, we enforce a unique decomposition algorithmically by first obtaining the warping functions through the pairwise comparisons, and then attributing the remaining variation to amplitude variation that is quantified in a second step. However, as outlined in Kneip and Ramsay (2008), while there are many registration procedures that will give rise to consistent registrations, the most meaningful criterion to determine whether observed variation is due to registration or amplitude variation is interpretability in the context of specific applications, which in our application is intrinsically linked to the nature of the relationship between the linguistic covariates and the functional principal component scores of both amplitude and warping functions. Emphasizing the linguistically important JND criteria, the eigenfunctions associated with the largest four eigenvalues in both the amplitude and phase bases could all be detected by the human ear, and as such, would affect the sound being perceived. Further, because we consider an LME model for the joint score vector associated with the amplitude and warping functions, we are able to capture correlations between the two sets of functions. This joint modeling helps alleviate some of the concerns regarding overall identifiability, as it is the joint rather than marginal results that are of interest. The fact that the scores and FPCs were all linguistically interpretable also gives further credence to the approach. Additionally, applying a different registration method (Zhang and Müller 2011) led to similar linguistic interpretations (see supplementary material, Section S.6).

In addition to the issue of identifiability, the obvious technical caveats with this work stem from three main areas: the discretization procedure, the time-registration procedure, and the multivariate mixed effects regression. Focusing on the discretization, the choice of basis is of fundamental importance. While we used principal components for the reason mentioned above, there have been questions as to whether a residual sum of squares optimality is most appropriate. It is certainly an open question when it comes to application specific cases (Bruns 2004). Aside from the case of parametric bases, nonparametric basis function generation procedures such as ICA (Hyvärinen and Oja 2000) have become recently increasingly more prominent. These bases could be used in the analysis, although the subsequent modeling of the scores would become inherently more complex due to the lack of certain orthogonality assumptions.

Regarding time-registration, there are a number of open questions regarding the choice of the framework to be used. However, we have examined two different frameworks and both these resulted in similar overall conclusions. The choice of the time-registration framework ultimately relies on the theoretical

assumptions one is willing to make and on the application and the samples to be registered. For the linguistic application, we are concerned with, it is not unreasonable to assume that the pairwise alignment corresponds well to the intuitive belief that intrinsically humans have a “reference” utterance onto which they “map” what they hear to comprehend it (Benesty, Sondhi, and Huang 2008).

Finally, multivariate mixed effects regression is itself an area with many possibilities. Optimization for such models is not always trivial and as the model and/or the sample size increases, estimation of the model tends to get computationally expensive. In our case, we used a hybrid optimization procedure that changes between a simplex algorithm (Nelder-Mead) and a quasi-Newton approach (Broyden-Fletcher-Goldfarb-Shanno; BFGS) (Kelley 1999) (see supplementary material for more information); in recent years research regarding the optimization tasks in an LME model has tended to focus on derivative free procedures. In a related issue, the choice of covariance structure is of importance. While we chose a very flexible covariance structure, the choice of covariance can convey important experimental insights. A final note specific to our problem was the presence of only five speakers. Speaker effect is prominent in many components and appears influential despite the small number of speakers available; nevertheless we recognize that including more speakers would have certainly been beneficial if they had been available. Given that the Speaker effect was the most important random-effect factor of this study, the inclusion of random slopes might also have been of interest (Schielzeth and Forstmeier 2009; Barr et al. 2013). Nevertheless, the inclusion of generic linear, quadratic, and cubic gender-specific down-drift effects presented through the break components allows substantial model flexibility to avoid potential design-driven misspecification of the random effects.

In conclusion, we have proposed a comprehensive modeling framework for the analysis of phonetic information in its original domain of collection, via the joint analysis of phase, amplitude, and duration information. The models are interpretable due to the LME structure, and estimable in a standard Euclidean domain via the compositional transform of the warping functions. The resulting model provides estimates and ultimately a typography of the shape, distortion and duration of tonal patterns, and effects in one of the world’s major languages.

SUPPLEMENTARY MATERIAL

Additional supplementary information is provided in the following files:

Supplementary Material Additional figures, information about covariance structures, and implementation details concerning the multivariate linear mixed effect analysis.

Supplementary Code Code and example analysis for multivariate linear mixed effect model.

[Received June 2015. Revised May 2015.]

REFERENCES

- Aitchison, J. (1982), “The Statistical Analysis of Compositional Data,” *Journal of the Royal Statistical Society*, Series B, 44, 139–177. [545,550]
- (1983), “Principal Component Analysis of Compositional Data,” *Biometrika*, 70, 57–65. [550,554]
- Aitchison, J., and Greenacre, M. (2002), “Biplots of Compositional Data,” *Journal of the Royal Statistical Society*, Series C, 51, 375–392. [550]
- Aston, J., Chiou, J., and Evans, J. (2010), “Linguistic Pitch Analysis using Functional Principal Component Mixed Effect Models,” *Journal of the Royal Statistical Society*, Series C, 59, 297–317. [545,547,548,551,553]
- Baayen, R., Davidson, D., and Bates, D. (2008), “Mixed-Effects Modeling With Crossed Random Effects for Subjects and Items,” *Journal of Memory and Language*, 59, 390–412. [546]
- Baayen, R. H. (2008), *Analyzing Linguistic Data. A Practical Introduction to Statistics using R*, Cambridge, UK: Cambridge University Press. [553]
- Barr, D. J., Levy, R., Scheepers, C., and Tily, H. J. (2013), “Random Effects Structure for Confirmatory Hypothesis Testing: Keep it Maximal,” *Journal of Memory and Language*, 68, 255–278. [558]
- Bates, D., Maechler, M., Bolker, B., and Walker, S. (2013), *lme4: Linear Mixed-Effects Models Using Eigen and S4, r package version 1.0-4*, Vienna, Austria: R Foundation for Statistical Computing. [546,552]
- Beckman, M. E., Hirschberg, J., and Shattuck-Hufnagel, S. (2006), “The Original ToBI system and the Evolution of the ToBI Framework,” in *Prosodic Typology: The Phonology of Intonation and Phrasing*, ed. S.-A. Jun, pp. 9–54. Oxford: Oxford University Press. [547]
- Benesty, J., Sondhi, M. M., and Huang, Y. (2008), *Springer Handbook of Speech Processing*, Berlin Heidelberg, Germany: Springer, Chapters 21, p. 445; 23, pp. 481–484; 36, pp. 731–732. [558]
- Benko, M., Härdle, W., and Kneip, A. (2009), “Common Functional Principal Components,” *The Annals of Statistics*, 37, 1–34. [548,552]
- Black, A. W., and Hunt, A. (1996), “Generating F0 Contours From to BI Labels Using Linear Regression,” in *ICSLP*, pp. 1385–1388. [549]
- Brumback, B. A., and Rice, J. A. (1998), “Smoothing Spline Models for the Analysis of Nested and Crossed Samples of Curves,” *Journal of the American Statistical Association*, 93, 961–976. [551]
- Bruns, A. (2004), “Fourier-, Hilbert- and Wavelet-Based Signal Analysis: Are They Really Different Approaches?” *Journal of Neuroscience Methods*, 137, 321–332. [557]
- Buser, P., and Imbert, M. (1992), *Audition* (1st ed.), Cambridge, MA: MIT Press, Chapter 2. [554]
- Cangelosi, R., and Goriely, A. (2007), “Component Retention in Principal Component Analysis With Application to cDNA Microarray Data,” *Biology Direct*, 2, 2+. [549]
- Castro, P., Lawton, W., and Sylvestre, E. (1986), “Principal Modes of Variation for Processes With Continuous Sample Curves,” *Technometrics*, 28, 329–337. [546]
- Central Intelligence Agency (2012), “The CIA World Factbook,” available at <https://www.cia.gov/library/publications/the-world-factbook/geos/xx.html>. [545]
- Chen, K., and Müller, H.-G. (2014), “Modeling Conditional Distributions for Functional Responses, With Application to Traffic Monitoring via GPS-Enabled Mobile Phones,” *Technometrics*, 56, 347–358. [548]
- Chiou, J., Müller, H., and Wang, J. (2003), “Functional Quasi-Likelihood Regression Models With Smooth Random Effects,” *Journal of the Royal Statistical Society*, Series B, 65, 405–423. [552]
- Egozcue, J. J., Pawlowsky-Glahn, V., Mateu-Figueras, G., and Barceló-Vidal, C. (2003), “Isometric Logratio Transformations for Compositional Data Analysis,” *Mathematical Geology*, 35, 279–300. [550]
- Evans, J., Chu, M., Aston, J., and Su, C. (2010), “Linguistic and Human Effects on F0 in a Tonal Dialect of Qiang,” *Phonetica*, 67, 82–99. [546]
- Filzmoser, P., Hron, K., and Reimann, C. (2009), “Principal Component Analysis for Compositional Data With Outliers,” *Environmetrics*, 20, 621–632. [550]
- Flury, B. N. (1984), “Common Principal Components in k Groups,” *Journal of the American Statistical Association*, 79, 892–898. [548]
- Fujisaki, H. (2004), “Information, Prosody, and Modeling-With Emphasis on Tonal Features of Speech,” in *Speech Prosody 2004, International Conference, ISCA*, pp. 1–10. [546,555,557]
- Grabe, E., Kochanski, G., and Coleman, J. (2007), “Connecting Intonation Labels to Mathematical Descriptions of Fundamental Frequency,” *Language and Speech*, 50, 281–310. [548]
- Greven, S., and Kneib, T. (2010), “On the Behaviour of Marginal and Conditional AIC in Linear Mixed Models,” *Biometrika*, 97, 773–789. [554]
- Gu, W., Hirose, K., and Fujisaki, H. (2006), “Modeling the Effects of Emphasis and Question on Fundamental Frequency Contours of Cantonese Utterances,” *IEEE Transactions on Audio, Speech, and Language Processing*, 14, 1155–1170. [545]
- Guo, W. (2002), “Functional Mixed Effects Models,” *Biometrics*, 58, 121–128. [552]

- Hadfield, J. D. (2010), "MCMC Methods for Multi-Response Generalized Linear Mixed Models: The MCMCglmm R Package," *Journal of Statistical Software*, 33, 1–22. [546,552]
- Hadjipantelis, P., Aston, J., and Evans, J. (2012), "Characterizing Fundamental Frequency in Mandarin: A Functional Principal Component Approach Utilizing Mixed Effect Models," *Journal of the Acoustical Society of America*, 131, 4651–4664. [551,552,554,555,557]
- Hirst, D., and Espesser, R. (1993), "Automatic Modelling of Fundamental Frequency Using a Quadratic Spline Function," *Travaux de l'Institut de phonétique d'Aix*, 15, 71–85. [557]
- Hyvärinen, A., and Oja, E. (2000), "Independent Component Analysis: Algorithms and Applications," *Neural Networks*, 13, 411–430. [557]
- Izenman, A. (2008), *Modern Multivariate Statistical Techniques: Regression, Classification and Manifold Learning*, New York: Springer Verlag, Chapter 6. [552]
- Jacewicz, E., Fox, R. A., and Wei, L. (2010), "Between-Speaker and Within-Speaker Variation in Speech Tempo of American English," *Journal of the Acoustical Society of America*, 128, 839–850. [549,554]
- Jurafsky, D., and Martin, J. H. (2009), *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition* (2nd ed.), Upper Saddle River, NJ: Prentice Hall PTR, Chapter 7.3. [545]
- Kelley, C. T. (1999), "Iterative Methods for Optimization," in *Frontiers in Applied Mathematics*, Philadelphia, PA: SIAM. [558]
- Kneip, A., and Gasser, T. (1992), "Statistical Tools to Analyze Data Representing a Sample of Curves," *Annals of Statistics*, 20, 1266–1305. [557]
- Kneip, A., and Ramsay, J. (2008), "Combining Registration and Fitting for Functional Model," *Journal of the American Statistical Association*, 103, 1155–1165. [545,557]
- Kochanski, G. (2006), "Prosody Beyond Fundamental Frequency," in *Methods in Empirical Prosody Research*, ed. S. Sudhoff, Berlin: Walter De Gruyter Inc, pp. 89–123. [549]
- Koening, L. L., Lucero, J. C., and Perlman, E. (2008), "Speech Production Variability in Fricatives of Children and Adults: Results of Functional Data Analysis," *Journal of the Acoustical Society of America*, 5, 3158–3170. [547]
- Kurtek, S., Srivastava, A., Klassen, E., and Ding, Z. (2012), "Statistical Modeling of Curves Using Shapes and Related Features," *Journal of the American Statistical Association*, 107, 1152–1165. [550]
- Laird, N. M., and Ware, J. H. (1982), "Random-Effects Models for Longitudinal Data," *Biometrics*, 38, 963–974. [545]
- Latsch, V., and Netto, S. L. (2011), "Pitch-Synchronous Time Alignment of Speech Signals for Prosody Transplantation," in *International Symposium on Circuits and Systems (ISCAS 2011)*, IEEE, pp. 2405–2408. [546]
- Lee, S., Byrd, D., and Krivokapic, J. (2006), "Functional Data Analysis of Prosodic Effects on Articulatory Timing," *Journal of the Acoustical Society of America*, 119, 1666–1671. [547]
- Leonard, T. (1973), "A Bayesian Method for Histograms," *Biometrika*, 60, 297–308. [546]
- Lindquist, M. A., Spicer, J., Asllani, I., and Wager, T. D. (2012), "Estimating and Testing Variance Components in a Multi-Level GLM," *Neuroimage*, 59, 490–501. [553]
- Liu, X., and Müller, H. G. (2004), "Functional Convex Averaging and Synchronization for Time-Warped Random Curves," *Journal of the American Statistical Association*, 99, 687–699. [557]
- Louw, J., and Barnard, E. (2004), "Automatic Intonation Modeling With INTSINT," in *Proceedings of the Pattern Recognition Association of South Africa*, pp. 107–111. [555]
- Mercer, J. (1909), "Functions of Positive and Negative Type, and Their Connection With the Theory of Integral Equations," *Philosophical Transactions of the Royal Society of London, Series A*, 209, 415–446. [549]
- Minka, T. P. (2001), "Automatic Choice of Dimensionality for PCA," *Advances in Neural Information Processing Systems*, 15, 598–604. [549]
- Mixdorff, H. (2000), "A Novel Approach to the Fully Automatic Extraction of Fujisaki Model Parameters," in *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing, 2000. ICASSP'00*, IEEE (Vol. 3), pp. 1281–1284. [557]
- Patterson, H. D., and Thomson, R. (1971), "Recovery of Inter-Block Information When Block Sizes are Unequal," *Biometrika*, 58, 545–554. [552]
- Pawlowsky-Glahn, V., and Egozcue, J. (2006), "Compositional Data and Their Analysis: An Introduction," *Geological Society, London, Special Publications*, 264, 1–10. [546]
- Pinheiro, J., Bates, D., DebRoy, S., Sarkar, D., and R Core Team (2013), *nlme: Linear and Nonlinear Mixed Effects Models, r package version 3.1-109*, Vienna, Austria: R Foundation for Statistical Computing. [552]
- Prom-On, S., Xu, Y., and Thipakorn, B. (2009), "Modeling Tone and Intonation in Mandarin and English as a Process of Target Approximation," *The Journal of the Acoustical Society of America*, 125, 405–424. [545]
- Quene, H. (2007), "On the Just Noticeable Difference for Tempo in Speech," *Journal of Phonetics*, 35, 353–362. [549,554]
- Rabiner, L. (1989), "A Tutorial on HMM and Selected Applications in Speech Recognition," *Proceedings of the IEEE*, 77, 257–286. [546]
- Ramsay, J., and Li, X. (1998), "Curve Registration," *Journal of the Royal Statistical Society, Series B*, 60, 351–363. [552]
- Ramsay, J., and Silverman, B. (2005), *Functional Data Analysis*, New York: Springer Verlag, Chapters 1, pp. 9–15; 3, pp. 38–45; 8, pp. 148–154. [545,547,548,552]
- Ramsay, J. O., Munhall, K., Gracco, V., and Ostry, D. (1996), "Functional Data Analyses of Lip Motion," *Journal of the Acoustical Society of America*, 99, 3718–3727. [547]
- Sakoe, H. (1979), "Two-Level DP-Matching—A Dynamic Programming-Based Pattern Matching Algorithm for Connected Word Recognition," *IEEE Transactions on Acoustics, Speech and Signal Processing*, 27, 588–595. [546]
- Schielzeth, H., and Forstmeier, W. (2009), "Conclusions Beyond Support: Overconfident Estimates in Mixed Models," *Behavioral Ecology*, 20, 416–420. [558]
- Searle, S. R., Casella, G., and McCulloch, C. E. (2006), *Variance Components* (2nd ed.), Hoboken, NJ: Wiley. [552]
- Su, Z., and Wang, Z. (2005), "An Approach to Affective-Tone Modeling for Mandarin," in *Affective Computing and Intelligent Interaction Lecture Notes in Computer Science*, Vol. 3784, eds. J. Tao, T. Tan, and R. Picard, Berlin-Heidelberg: Springer, pp. 390–396. [545]
- Tang, P., and Müller, H. (2008), "Pairwise Curve Synchronization for High-Dimensional Data," *Biometrika*, 95, 875–889. [548,549,550,552,557]
- (2009), "Time-Synchronized Clustering of Gene Expression Trajectories," *Biostatistics*, 10, 32–45. [547,548,549]
- Taylor, P. (2000), "Analysis and Synthesis of Intonation Using the Tilt Model," *The Journal of the Acoustical Society of America*, 107, 1697–1714. [555]
- Torgerson, R. C. (2005), "A Comparison of Beijing and Taiwan Mandarin Tone Register: An Acoustic Analysis of Three Native Speech Styles," M.Sc. dissertation, Brigham Young University, Provo, UT. [553]
- Tseng, C., Cheng, Y., and Chang, C. (2005), "Sinica COSPRO and Toolkit: Corpora and Platform of Mandarin Chinese Fluent Speech," in *Proceedings of Oriental COCODSA*, pp. 6–8. [547]
- Tucker, J. D., Wu, W., and Srivastava, A. (2013), "Generative Models for Functional Data Using Phase and Amplitude Separation," *Computational Statistics & Data Analysis*, 61, 50–66. [550]
- Wilkinson, G. N. and Rogers, C. E. (1973), "Symbolic Description of Factorial Models for Analysis of Variance," *Journal of the Royal Statistical Society, Series C*, 22, 392–399. [552]
- Xu, Y. (1999), "Effects of Tone and Focus on the Formation and Alignment of f0 Contours," *Journal of Phonetics*, 27, 55–105. [553]
- Xu, Y., and Wang, Q. E. (2001), "Pitch Targets and Their Realization: Evidence From Mandarin Chinese," *Speech Communication*, 33, 319–337. [545]
- Yoshioka, T., Sehr, A., Delcroix, M., Kinoshita, K., Maas, R., Nakatani, T., and Kellermann, W. (2012), "Making Machines Understand us in Reverberant Rooms: Robustness Against Reverberation for Automatic Speech Recognition," *IEEE Signal Processing Magazine*, 29, 114–126. [546]
- Zhang, Z., and Müller, H.-G. (2011), "Functional Density Synchronization," *Computational Statistics and Data Analysis*, 55, 2234–2249. [550,557]